

Analog Forecasting with Dynamics-Adapted Kernels

Zhizhen Zhao

Department of Mathematics and Center for Atmosphere and Ocean Science, Courant
Institute of Mathematical Sciences, New York University, New York, NY, USA

E-mail: jzhao@cims.nyu.edu

Dimitrios Giannakis

Department of Mathematics and Center for Atmosphere and Ocean Science, Courant
Institute of Mathematical Sciences, New York University, New York, NY, USA

E-mail: dimitris@cims.nyu.edu

Abstract. Analog forecasting is a nonparametric technique introduced by Lorenz in 1969 which predicts the evolution of states of a dynamical system (or observables defined on the states) by following the evolution of the sample in a historical record of observations which most closely resembles the current initial data. Here, we introduce a suite of forecasting methods which improve traditional analog forecasting by combining ideas from kernel methods developed in harmonic analysis and machine learning and state-space reconstruction for dynamical systems. A key ingredient of our approach is to replace single-analog forecasting with weighted ensembles of analogs constructed using local similarity kernels. The kernels used here employ a number of dynamics-dependent features designed to improve forecast skill, including Takens' delay-coordinate maps (to recover information in the initial data lost through partial observations) and a directional dependence on the dynamical vector field generating the data. Mathematically, our approach is closely related to kernel methods for out-of-sample extension of functions, and we discuss alternative strategies based on the Nyström method and the multiscale Laplacian pyramids technique. We illustrate these techniques in applications to forecasting in a low-order deterministic model for atmospheric dynamics with chaotic metastability, and interannual-scale forecasting in the North Pacific sector of a comprehensive climate model. We find that forecasts based on kernel-weighted ensembles have significantly higher skill than the conventional approach following a single analog.

Keywords: Analog forecasting, kernel methods, out-of-sample extension , delay-coordinate maps

1. Introduction

In many branches of science and engineering, advances in modeling capabilities and the proliferation of experimental and observational data have spurred the development of new mathematical and computational techniques for predicting the dynamical behavior of complex nonlinear systems with imperfect models and partial knowledge of the initial data. The Earth’s climate system is a classical example of such a complex and partially observed system coupling physical processes from the atmosphere, ocean, and land over a wide range of spatial and temporal scales. Currently, as well as in the foreseeable future [29], the equations of motion for the full climate system cannot be integrated via direct numerical simulation and must be parameterized to account for unresolved scales, while in some cases the equations of motion are only partially known [11]. A topic of significant current interest is how to make accurate and reliable predictions of large-scale climate patterns over timescales ranging from weeks to decades, using imperfect models initialized with the incomplete information about the current state of the system provided by observational networks [48, 44, 21].

One way of making initial-value forecasts in complex systems is through large-scale numerical simulation. For example, in climate science, regional weather models and general circulation models (GCMs) are used extensively to make forecasts over lead times spanning a few hours (in weather forecasting) to several decades (in long-term climate projections). Broadly speaking, these numerical models are based on a dynamical core for the atmosphere and/or ocean derived from the Navier-Stokes equations on a rotating sphere, which is coupled to a complex network of physics models describing processes such as radiative energy transfer and phase transitions. In an operational environment, the models are coupled with a data assimilation system providing initial conditions from the (incomplete) data available from sensors such as satellites, weather stations, and buoys. The significant increases in forecast skill for the weather and climate over the past decades is due to advances in both the forward models and data assimilation systems [48].

Distinct from large-scale numerical models are parametric and nonparametric empirical low-order methods [36, 22]. These techniques do not rely on a large-scale forward model in the forecast step. Rather, predictions are made using a model trained on a reference time series (which may be obtained from historical observations of nature, or a large-scale numerical model), and initialized with the currently observed data from nature. The empirical models generally make coarser forecasts and may also underperform large-scale models in terms of innovation, but their strengths lie in their low computational cost and ability to be trained using data from nature. Low computational cost is desirable in scenarios where real-time or near real-time forecasts are required but the cost of frequent large-scale model runs is prohibitive [64], and in applications where low-order process models are embedded in large-scale models as subgrid-scale parameterization schemes [24]. Equally important, low-order empirical models can be trained using data acquired from nature, therefore avoiding some of the dynamical model errors present in large-scale forward models. For instance, a major barrier in extending the skill of numerical weather forecasts at timescales greater than about

a week is the poor dynamical representation of organized tropical convection [45]. As a result, low-order empirical models frequently outperform their large-scale counterparts in predicting climate phenomena which are the outcome of convective organization, including the Indian monsoon [65] and the Madden-Julian oscillation [40].

Within empirical low-order models, a dichotomy can be made between parametric and nonparametric techniques (though the dichotomy is loose, as there exist semiparametric techniques combining aspects of both approaches; e.g., [9]). Parametric techniques require the specification of an explicit model for the dynamics, whose parameters are estimated statistically from past observations of the system. Classical examples of such techniques are linear and nonlinear regression models [36] and Bayesian hierarchical models [22]. When available, prior knowledge of physical laws, or physics-based constraints [42], can significantly facilitate the parametric model design with a corresponding improvement of skill (e.g., [15]). Yet, in certain applications, prior knowledge of an appropriate parametric model is not available. Such an application, which we will study in Section 6.2 ahead, is the low-frequency variability of the North Pacific Ocean. There, even state-of-art parametric models generally provide poorer forecast skill than simple persistence forecasts [20]. In such scenarios, nonparametric models are a viable alternative strategy relying on only general conditions on the properties of the dynamics (e.g., ergodicity).

Nonparametric techniques utilize past observations of a dynamical system to infer its future behavior without assuming an explicit parametric form of the equations of motion. As a result, their main advantage is their applicability in a wide range of forecasting scenarios. At the same time, nonparametric models must compensate for their flexible structure through plentiful samples in phase space (whereas in parametric techniques a partial sampling of the phase space may be sufficient for accurate parameter estimation). Classical nonparametric forecasting techniques include local linear models [27, 51, 39] and analog forecasting [41, 62]; the methods developed in this paper can be viewed as a generalization of the latter technique using ideas from kernel methods in harmonic analysis and machine learning [17, 10, 18, 16, 28, 50]. The methods developed here also have connections with a recently developed nonparametric method called diffusion forecasting [7] that makes use of related kernel techniques, and we will discuss these connections in Section 5.

1.1. Analog forecasting

The focus of this paper is on nonparametric methods inspired by the analog forecasting technique. This technique was introduced by Lorenz in 1969 [41] as a method for predicting the time evolution of observables in dynamical systems based on a historical record of training data. In the initialization stage of analog forecasting, one identifies an analog, i.e., the state in the historical record which most closely resembles the current initial data. Then, in the forecast step, the historical evolution of that state is followed for the desired lead time, and the observable of interest is predicted based on its value on the analog. Unlike prediction using GCMs and parametric low-order models, the analog approach makes no

assumptions about the dynamics of the system being examined. Thus, if the historical record of observations comes from nature, the method avoids the dynamical model errors of parametric models essentially by construction, allowing the deployment of the method to real-world problems where parametric models are known to perform poorly [60, 65]. Analog forecasting is also useful in situations where the reference time series itself is the output of numerical models. For instance, in [12] analogs are used to assess the long-range internal predictability properties of GCMs. Nonlinear forecasting based on libraries of past histories has also been used in ecology [55].

Three major factors influencing the efficacy of analog forecasting are (i) the identification of “skillful” analogs in the training data; (ii) the choice of forecast observable (predictand); (iii) model error in the training data. Geometrically, the ability to identify skillful analogs amounts to being able to identify subsets of the training data whose dynamical evolution will shadow the future time evolution for the given initial data. In general, the required shadowing accuracy to achieve a given forecast accuracy depends on both the dynamical system and the observable, as different observables have different predictability limits. For instance, the predictability of the full state vector is limited by the maximal Lyapunov exponent of the dynamical system, and a trivial observable equal to a constant is predictable at arbitrarily long lead times. Indeed, an important problem in data-driven forecasting techniques is to identify observables which are physically meaningful but also predictable. The properties of the prediction observable also affects the reliability of analog forecasts in situations where the training data have model error. In particular, in some situations it may be preferable (e.g., when the time span of the available data from nature is short), or necessary (e.g., when the predictand is not directly observable) to use a training dataset generated by an imperfect model, despite that analog forecasting cleverly avoids dynamical model error if the training data are acquired for nature.

In real-world applications, the problem of analog identification is compounded by the fact that the initial data are often high-dimensional and incomplete (i.e., insufficient to uniquely determine the future dynamical evolution of the system). Moreover, highly predictable observables are not known a priori. As a concrete example, consider climate forecasting in the North Pacific Ocean given observations of sea surface temperature (SST) over a spatial grid. In this example, which will be studied in Section 6.2, the dimension of the initial data is equal to the number of spatial gridpoints where SST measurements are taken. For the spatial resolutions typical of current-generation observational networks and GCMs ($\lesssim 1^\circ$) the number of North Pacific gridpoints is $O(10^4)$. However, knowledge of SST at a single instance of time is not sufficient to determine its future evolution, as many of the essential degrees of freedom of the atmosphere and the ocean (e.g., the vertically-resolved circulation, temperature, and density fields) are not represented in SST snapshots. As a result, similar SST snapshots may be produced by states lying far apart on the attractor, and identifying analogs solely on the basis of similarity of SST snapshots is likely to lead to poor prediction skill.

Regarding the choice of prediction observable, considerable efforts are being made by oceanographers to identify patterns which characterize the variability of the ocean on timescales of several years. In the North Pacific, one such prominent patterns are the Pacific decadal oscillation (PDO) [43], which is conventionally obtained by applying principal components analysis (PCA) to seasonally-detrended and low-pass filtered SST data. In general, an observable measuring the PDO activity [e.g., the principal component (PC) corresponding to the pattern] is subjective, in the sense that it is defined through a data analysis algorithm applied to a finite dataset. PCA-type algorithms, in particular, are known to have poor skill in extracting dynamically significant patterns in certain nonlinear dynamical systems [2, 23], motivating the development of alternative data-driven approaches for extracting observables for empirical forecasting.

1.2. Contributions of this work

The central theme of this work is that kernel methods, in conjunction with state-space reconstruction methods for dynamical systems, can lead to significant improvements in the traditional formulation of analog forecasting, benefiting both the identification of skillful analogs, as well as the construction of intrinsic slow modes for prediction. To motivate our approach, we note that a generic feature of many dynamical systems of interest in science and engineering is that, after initial transients have decayed, the system evolves on low-dimensional subsets of phase space (attractors) [37, 25]. These nonlinear structures are embedded in the high-dimensional ambient data space in potentially complicated ways, meaning that the best-fitting linear hyperplane is likely to be of significantly higher dimension than the intrinsic dimension of the data. Kernel algorithms [3, 17, 18, 59, 50, 10, 28] are well-suited to handle nonlinear geometrical structures of data such as manifolds. Here, our objective is to improve the traditional analog forecasting using these techniques, adapted to deal with data generated by dynamical systems [32, 6, 30].

For our purposes, a kernel will be a measure of similarity between pairs of samples in data space. Suitably normalized [17, 10], these objects naturally lead to local averaging operators acting on observables such that. In particular, given previously unseen initial data, these operators act by averaging the observable values over the samples of the training data which are most similar to the initial data with respect to the kernel. Conceptually, averaging extends the notion of single-analog identification in conventional analog forecasting to the construction of a weighted ensemble of analogs with weights determined by the averaging operator.

For suitably constructed kernels, the behavior of the associated averaging operators can be understood through a Riemannian metric on the nonlinear manifold sampled by the data that depends on the kernel [10]. In this context, kernels can therefore be thought of as inducing a Riemannian geometry, sometimes referred to as diffusion geometry, to the training data. Having in mind that the skill of analog forecasting depends on the ability to geometrically identify good analogs in the training data, we propose to use the properties

of a geometry adapted to the dynamical system generating the data to improve the efficacy of this step. In particular, following [32, 33, 6, 30], our approach is to take advantage of a special structure present in data generated by dynamical systems, namely the time ordering of the samples which is the outcome of the dynamics.

The time-ordering of the data samples has been utilized at least since the early 1980s in state-space reconstruction methods [57, 13, 52]. In these works, it was established that under generic conditions one can recover the topology of the attractor of a dynamical system from time-ordered partial observations by embedding the data into a higher-dimensional space consisting of lagged sequences of observations over a temporal window. At the very least, when the initial data are incomplete, this suggests to perform analog forecasting using Euclidean distances or similarity kernels defined in delay-coordinate space. Besides topological aspects, however, delay-coordinate mapping also influences the geometry of the data irrespective of whether or not the initial data is complete. That is, pairwise distances on the data manifold in delay coordinate space depend not only on “snapshots” at single time instances, but on dynamical trajectories (“videos”). In [6], it was shown that adding delays increasingly biases the kernel-induced metric towards the most stable subspace of the Lyapunov metric, enhancing the kernel’s timescale separation capability. This approach was also used heuristically in earlier work on so-called nonlinear Laplacian spectral analysis (NLSA) algorithms [32, 31, 33], where kernels in delay-coordinate space were found to have significantly higher skill in recovering patterns analogous to the PDO in North Pacific data from a GCM. Moreover, a set of spatial patterns recovered via NLSA were found to yield Galerkin-reduced models reproducing faithfully the chaotic regime transitions in a low-order model for the atmosphere [23] where PCA is known to fail dramatically. The low-order atmospheric model and the high-dimensional North Pacific GCM data will both be employed here as applications of our kernel analog forecasting techniques. Moreover, we will study an example with model error in the training data simulated by treating North Pacific SST output from a GCM operating at a given atmospheric resolution as “nature” and the corresponding output obtained from a model with a coarser atmosphere as training data from an imperfect model.

Besides delay-coordinate mappings, the kernel used in NLSA features an additional dependence on the time-tendency of the data evaluated through finite differences of the data in time. In [30], it was shown that the finite-differenced data provide an empirical approximation of the dynamical vector field on the attractor generating the dynamics. That work also established that by incorporating a suitable directional dependence, biasing the kernel towards states which are mutually aligned with the dynamical flow, one gains additional invariance properties and ability to extract slow dynamical time scales. These so-called cone kernels will be employed here as our preferred similarity kernel to identify skillful analogs in the training data, and we will also use kernel eigenfunctions associated with cone kernels to construct slow observables with high predictability.

With any kernel, the ability to perform analog forecasts hinges on the ability to

evaluate observables at points lying outside the training dataset. Here, we carry out these tasks using two popular out-of-sample extension techniques, namely the Nyström method [47, 5, 18] and Laplacian pyramids [50, 28]. The Nyström method performs particularly well when the prediction observable is itself a kernel eigenfunction (which will be the case in the North Pacific SST application), but may become ill-conditioned and/or biased for more general observables. In those situations, the multiscale procedure for out-of-sample extension in Laplacian pyramids is preferable. Used in conjunction with a bistochastic kernel normalization [16], Laplacian pyramids is also well-suited for analog forecasting with incomplete initial data; a situation encountered frequently in real-world applications. Both the Nyström method and Laplacian pyramids perform out-of-sample extension through one or more kernel integral operators acting on observables of the dynamical system. Thus, analog forecasting with these techniques can be thought of as an ensemble generalization of conventional single-analog forecasting. In applications, we find that these methods have significantly higher forecast skill than what is possible with single analogs.

This paper is organized as follows. In Section 2, we describe the mathematical framework of kernel analog forecasting, focusing on the properties of analog forecasting with Nyström extension and Laplacian pyramids for general types of kernels. We discuss specific examples of dynamics-adapted kernels and their normalization in Section 3. In Section 4, we discuss forecast error metrics and the uncertainty quantification of our scheme. In Section 5, we provide a comparison of kernel analog forecasting with the diffusion forecast method [7]. In Section 6, we present applications of kernel analog forecasting in a low-order model for the atmosphere and the North Pacific sector of comprehensive climate models. We conclude in Section 7.

2. Kernel analog forecasting

In this Section, we present the mathematical framework of our kernel analog forecasting strategies. These techniques are inspired by Lorenz’s analog method [41], so, after establishing our notation and main assumptions in Section 2.1, we provide an overview of conventional forecasting in Section 2.2. In Section 2.3, we introduce a basic version of kernel analog forecasting based on averaging with a single operator. While this would not be our method of choice in practice, it serves as an illustration of the main ideas and challenges in kernel analog forecasting. In Sections 2.4 and 2.5, we present refinements of the basic kernel analog forecasting approach using Nyström out-of-sample extension [18] and Laplacian pyramids [50, 28].

2.1. Notation and preliminaries

Consider time-ordered samples $\{x_0, x_1, \dots, x_{N-1}\}$ of a signal in a space X taken uniformly at times $\{t_0, t_1, \dots, t_{N-1}\}$ with $t_i = (i-1)\tau$. The signal is generated by a dynamical system, and observations of $y \in X$ at a single instance of time or over a finite time interval will serve

as initial data to forecast a quantity of interest (a prediction observable) f . The prediction observable is generated by the same dynamical system generating $\{x_i\}$, takes values in a space B , and the values $\{f_i\}$ at the sampling times $\{t_i\}$ are also observed.

Mathematically, we consider that the dynamical system evolves in an abstract (unobserved) phase space A and the observed data in X are the outcome of a mapping $\pi : A \mapsto X$ sending the system state $a_i \in A$ at time t_i to the data sample $x_i = \pi(a_i) \in X$. Similarly, the prediction observable is a map $f : A \mapsto B$ such that $f_i = f(a_i)$. For instance, in the ocean application of Section 6.2, we have $X = \mathbb{R}^d$ and $B = \mathbb{R}$, where $x_i \in X$ is the SST field measured at d spatial gridpoints in the North Pacific ocean at time t_i , and $f_i \in B$ is the corresponding value of a low-frequency mode of variability such as the PDO. Moreover, A is the phase space of the Earth's climate system. In general, we are interested in a partial-observations scenario where x_i provides incomplete knowledge of the system state a_i , i.e., π is a non-invertible map on its image. We call such initial data incomplete. On the other hand, we say that the initial data are complete if π is invertible on its image. In the applications of Section 6.2, observations of North Pacific SST alone are not sufficient to uniquely determine the full state of the climate system and the future evolution of the PDO, hence the initial data are incomplete.

Throughout, we assume that A and X are measurable spaces with σ -algebras \mathcal{A} and \mathcal{X} , respectively, and π is a measurable map from (A, \mathcal{A}) to (X, \mathcal{X}) . In the case of X , we also assume that it is a metrizable space (though we will not explicitly use its topology), and that \mathcal{X} is its Borel σ -algebra. We assume that the image $M = \pi(A)$ of the phase space under π is measurable, and denote the σ -algebra $\mathcal{X} \cap M$ by \mathcal{M} . Also, we denote the finite set of observations $\{x_i\} \subseteq M$ by \hat{M} .

Turning to the prediction observable, we consider that B is a Banach space with the corresponding Borel σ -algebra \mathcal{B} and that $f : (A, \mathcal{A}) \mapsto (B, \mathcal{B})$ is measurable. In what follows, we develop our methodology for scalar-valued observables, $B = \mathbb{R}$, but a similar treatment can be applied for more general classes of vector- and function-valued observables for which ergodic theorems apply [38]. In general, we distinguish between prediction observables defined objectively for the system, and observables constructed through a data analysis technique. In the low-order atmospheric model studied in Section 6.1, f will be the components of the state vector that capture metastable regime transitions—these are examples of objectively-defined observables. On the other hand, in Section 6.2 the goal will be to predict the time-evolution of the top two low-frequency modes constructed by processing the data via kernel algorithms. Desirable features of such data-driven observables is to evolve on slow intrinsic dynamical timescales, and have high smoothness as functions on the phase space A . In Section 3 ahead, we will see that dynamics-adapted kernels provide a promising route for constructing physically meaningful data-driven observables with favorable predictability properties.

In this work, we consider that the dynamics on (A, \mathcal{A}) is governed by a group or semigroup of measure-preserving transformations $\Phi_t : A \mapsto A$ with time parameter t and

invariant measure α ; that is, $\alpha(\Phi_t^{-1}(S)) = \alpha(S)$ for all $S \in \mathcal{A}$. The invariant measure of the dynamics induces a probability measure μ on (M, \mathcal{M}) such that $\mu(S) = \alpha(\pi^{-1}(S))$ for $S \in \mathcal{M}$. We also assume that Φ_t is ergodic, i.e., $\alpha(S)$ is equal to either 0 or 1 for every Φ_t -invariant set S . Let T be the time domain of the system. For concreteness, we develop our approach in the continuous-time invertible (group) case, $T = \mathbb{R}$, $\Phi_s \circ \Phi_t = \Phi_{s+t}$, $\Phi_t^{-1} = \Phi_{-t}$, and we assume that the dynamical system sampled at the discrete times $\{\dots, -2\tau, -\tau, 0, \tau, 2\tau, \dots\}$ is also ergodic. Counterparts of our results for continuous-time sampling follow by replacing discrete time averages $N^{-1} \sum_{i=0}^{N-1} f_i$ by continuous time integrals $t^{-1} \int_0^t f(t') dt'$. Our methods are also applicable in non-invertible (semigroup) systems, $T = \{t \geq 0\}$, with the exception of constructions utilizing backwards delay-coordinate maps (see Section 3.1).

At a minimum, we require that the prediction observable f is an integrable scalar-valued function in $L^1(A, \alpha)$, i.e.,

$$\bar{f} = \int_A f d\alpha < \infty. \quad (2.1)$$

This requirement will be sufficient for the most basic formulation of our analog forecasting techniques (see Section 2.3). In several other instances (e.g., the method based on Nyström extension in Section 2.4), we additionally require that f is a square integrable-observable in $L^2(A, \alpha) \subset L^1(A, \alpha)$. We denote the inner product and norm of $L^2(A, \alpha)$ by $\langle f, g \rangle = \int_A f g d\alpha$ and $\|f\| = \langle f, f \rangle^{1/2}$, respectively.

By the Birkhoff pointwise ergodic theorem, the time average of an observable f in $L^1(A, \alpha)$ along the sampled states a_i ,

$$\hat{f} = \frac{1}{N} \sum_{i=0}^{N-1} f_i, \quad f_i = f(a_i), \quad (2.2)$$

converges as $N \rightarrow \infty$ to \bar{f} for α -almost any starting point a_0 . Intuitively, we think of $N \rightarrow \infty$ as the limit of large data. The time average in (2.2) can also be expressed as an expectation value with respect to the purely atomic probability measure $\hat{\alpha} = N^{-1} \sum_{i=0}^{N-1} \delta_{a_i}$, where the δ_{a_i} are Dirac measures on (A, \mathcal{A}) centered at a_i . That is, we have $\hat{f} = \int_A f d\hat{\alpha}$. The ability to approximate expectation values such as (2.1) with the corresponding empirical time averages in (2.2) is crucial for the analog forecasting techniques developed below.

2.2. Conventional analog forecasting

Broadly speaking, analog forecasting [41] produces a forecast $\hat{F}_t(y)$ of an observable f of a dynamical system at forecast lead time t given initial data y in the space X . Assuming that t is a non-negative integer multiple k of the sampling interval τ , the necessary ingredients needed to carry out this task are (i) the training time series $\hat{M} = \{x_i\}$ with $x_i = \pi(a_i)$ lying in the same space as y ; (ii) the corresponding time series of the observable values $f_i = f(a_i)$; (iii) a distance function $D : X \times X \mapsto \mathbb{R}$. Typically, when $X = \mathbb{R}^n$, D is set to the Euclidean distance $D(y, x) = \|y - x\|$. In a perfect-model scenario, the training samples (x_i, f_i) are

generated by the same dynamical system as the dynamical system being forecast—this would be the case if x_i are observations of nature. Mathematically, this means that the initial data are the outcome of the same observation map as the training data, i.e., we have

$$y = \pi(b), \quad b \in A. \quad (2.3)$$

In an imperfect-model scenario, the (x_i, f_i) have systematic biases compared to the time evolution of these quantities in nature. Moreover, there does not necessarily exist a state b in the phase space of the imperfect model underlying the data y at forecast initialization as in (2.3). Such situations arise when the training data have been generated by a numerical model with dynamical errors relative to nature. However, as stated in Section 1.1, it may be desirable to use a long integration from a numerical model to improve the sampling density of state space despite model errors. Moreover, in some cases, the pairs (x_i, f_i) may not be observationally accessible. A notable example from climate science where this issue is encountered is forecasting and estimation of the interior circulation and thermal structure of the ocean. These variables are not accessible via remote sensing from satellites, and are only sparsely known from in situ sensors such as gliders and floats. For the remainder of this section we will restrict attention to the perfect model scenario where (2.3) holds.

Let U_t , $t \in T$, be the operator on $L^1(A, \alpha)$ whose action is given by composition with the evolution map:

$$U_t f = f \circ \Phi_t, \quad U_t f(a) = f(\Phi_t(a)). \quad (2.4)$$

That is, $U_t f$ is a time-shifted observable whose value at a is given by evaluating f at the point $\Phi_t(a)$ lying in the future of a at time t . Intuitively, $f_t(a)$ is the value of a “perfect” forecast for f at lead time t initialized at the dynamical state $a \in A$. Note that for $t = k\tau$ sufficiently small, $k \in \mathbb{N}$, and $a_i \in \{a_0, \dots, a_{N-1}\}$, $U_t f(a_i) = f_{i+k}$ is given by a k -step forward shift of the $\{f_i\}$ time series. In ergodic theory, the operators U_t are known as Koopman operators [26]. The set $\{U_t\}$ forms a group (or a semigroup in the non-invertible case) under composition of operators with $U_s \circ U_t = U_{s+t}$ and (in the group case) $U_t^{-1} = U_{-t}$. Koopman operators are similarly defined for observables in the Hilbert space $L^2(A, \alpha) \subset L^1(A, \alpha)$. These operators are isometric, $\|U_t f\| = \|f\|$, and are unitary with $U_t^* = U_t^{-1} = U_{-t}$ if $\{U_t\}$ is a group.

Analog forecasting is a two-step procedure which involves finding the sample x_i in the training data lying closest to y with respect to the distance on X , and then time-advancing the observable values f_i to produce a forecast. That is, one first identifies the analog $a_i \in A$ underlying the initial data y via

$$i = \operatorname{argmin}_{j \in \{1, \dots, n\}} D(y, x_j), \quad x_j = \pi(a_i), \quad (2.5)$$

and then evaluates the forecast at lead time $t = k\tau$ via

$$\hat{F}_t(y) := f_{i+k} = U_t f(a_i). \quad (2.6)$$

Note that the process of analog identification is implicit in the sense that in order to evaluate (2.6) it suffices to know the timestamp i of a_i from (2.5), but explicit knowledge

of the state $a_i \in A$ is not needed. Moreover, in (2.5) we have tacitly assumed that the minimizer of $D(y, \cdot)$ is unique in the training dataset \hat{M} ; if multiple minimizers are found (an unlikely situation in practice), then (2.6) can be replaced by the average forecast over those minimizers.

Conventional analog forecasting also has a measure-theoretic interpretation which will be useful in the development of kernel analog forecasting ahead. In particular, $\hat{F}_t(y)$ is equal to the expectation value of $U_t f$ with respect to the Dirac measure centered at the analog, i.e.,

$$\hat{F}_t(y) = \int_A U_t f d\delta_{a_i}. \quad (2.7)$$

Given the training data \hat{M} and the initial data y from (2.3), the forecast error (residual) of conventional analog forecasting is

$$\hat{r}_t(b) = U_t f(b) - \hat{F}_t(y) = U_t f(b) - U_t f(a_i) = \int_A U_t f(d\delta_b - d\delta_{a_i}). \quad (2.8)$$

Note that in an imperfect-model scenario $\hat{F}_t(y)$ would generally be evaluated for initial data lying outside of M , and (2.3) and (2.8) would not hold.

It is evident from (2.5)–(2.7) that analog forecasting is a fully nonparametric approach, relying solely on the geometric identification of nearest neighbors in the training data and empirical time shifts of time series. These properties make the method especially attractive in situations where one does not have knowledge of the equations of motion, or an appropriate parametric model is not available. For instance, in the ocean application in Section 6.2 construction of parametric regression models for the low-frequency modes is especially hard, and linear autoregressive models have worse skill than a trivial persistence forecast [20].

Yet, despite its attractive features, conventional analog forecasting suffers from high risk of overfitting the training data. As is evident from (2.6), the conventional analog forecast (and hence the corresponding forecast error in (2.8)) is a highly non-smooth function of the initial data—as y varies on X , $\hat{F}_t(y)$ changes discontinuously to the value $U_t f(a_i)$ of the observable on the analog that happens to minimize the distance to y , and the forecast error can exhibit large variations under small perturbations of the initial data. Moreover, it follows from (2.7) that conventional analog forecasting utilizes information from a single state in the training data to predict the future of the observable, but one would expect that other states in the training data also carry useful predictive information. Our objective in Sections 2.3–2.5 ahead is to address these shortcomings using kernel methods.

2.3. Basic kernel analog forecasting

According to the discussion in Section 2.2, conventional analog forecasting can be described as a mapping $y \mapsto \delta_{a_i}$ of the initial data in X to a Dirac measure δ_{a_i} on (A, \mathcal{A}) , followed by computation of the expectation value of the time-shifted observable $U_t f$ with respect to δ_{a_i} . In the simplest version of our kernel analog forecasting techniques, we replace δ_{a_i}

by a probability measure $\hat{\nu}_y$ with the properties that (i) expectation values with respect to $\hat{\nu}_y$ depend continuously on y ; (ii) in the limit of large data, $N \rightarrow \infty$, $\hat{\nu}_y$ converges to a probability measure ν_y which is absolutely continuous with respect to the invariant measure α of the dynamics.

A key ingredient of kernel analog forecasting is a kernel function, $K : X \times X \mapsto \mathbb{R}$, mapping pairs of points in the observation space X to a non-negative real number. Intuitively, we think of $K(y, x)$ as a measure of similarity between the data samples x and y , in the sense that large values of $K(y, x)$ indicate that x and y are highly similar, and values of $K(y, x)$ close to zero indicate that x and y are highly dissimilar. In what follows, x will typically lie in the training dataset \hat{M} whereas y will be a previously unseen sample. Mathematically, we require that K has the following properties:

- (i) K is non-negative;
- (ii) The functions $K(y, \cdot)$ and $K(\cdot, x)$ are measurable for every $y \in X$ and $x \in M$;
- (iii) There exists a constant $C < \infty$ such that $K \leq C$;
- (iv) There exists a constant $c > 0$ such that $\int_M K(y, \cdot) d\mu \geq c$ for all $y \in X$.
- (v) The function $K(\cdot, x)$ is continuous for every $x \in M$.

Among these requirements, properties (iii) and (iv) are not essential, but are very convenient for the exposition below. These assumptions could be weakened so long as quantities derived from K (e.g., the density ρ ahead) remain sufficiently smooth. Property (v) is also non-essential and can be relaxed if a continuous dependence of the kernel analog forecast on the initial data is not required.

By properties (ii) and (iii) and the fact that μ is finite, $K(y, \cdot)$ and $K(\cdot, x)$ are functions in $L^p(M, \mu)$, $1 \leq p \leq \infty$. Similarly, $K(y, \pi(\cdot))$ and $K(\pi(\cdot), x)$ are in $L^p(A, \alpha)$. In particular, the ergodic theorem in (2.2) applies, and the time average $\hat{q}(y)$, where

$$\hat{q}(y) = \int_M K(y, \cdot) d\hat{\mu} = \int_A K(y, \pi(\cdot)) d\hat{\alpha} = \frac{1}{N} \sum_{i=0}^{N-1} K(y, x_i), \quad (2.9)$$

converges to

$$q(y) = \int_M K(y, \cdot) d\mu = \int_A K(y, \pi(\cdot)) d\alpha.$$

By property (iii), $q(y)$ is bounded away from zero, and therefore for α -almost every starting point $a_0 \in A$ and for sufficiently large N , $\hat{q}(y)$ is nonzero—hereafter, we will assume that this is the case. The quantities $\hat{q}(y)$ and $q(y)$ will be useful in various kernel normalizations throughout the paper.

A basic example of a kernel that meets the properties listed above in the case that X is equipped with distance function D is the radial Gaussian kernel,

$$K_\epsilon(y, x) = e^{-D^2(y, x)/\epsilon},$$

where ϵ is a positive bandwidth parameter. This kernel is continuous with respect to both of its arguments. Radial Gaussian kernels are useful for approximating heat kernels on manifolds, and have been used extensively in harmonic analysis and machine learning [3, 17, 35, 59, 17, 53]. Note that if M and/or X are noncompact, the radial Gaussian kernel may fail to meet property (iv). Such situations can be handled by replacing (2.3) via so-called variable-bandwidth kernels [8], where the squared distance in the exponential function is scaled by a bandwidth function becoming large at distant points with small sampling density. In Section 3.2, we will discuss anisotropic Gaussian kernels which also take certain properties of the dynamical system generating the data into account.

To construct a probability measure on (A, \mathcal{A}) , we normalize the kernel and obtain weights $\hat{\rho}(y, x_i)$ such that $\hat{\rho}(y, x_i) \geq 0$ and $N^{-1} \sum_{i=0}^{N-1} \hat{\rho}(y, x_i) = 1$. There are several ways to perform this normalization [17, 16, 10], but arguably the simplest one is to divide $K(y, x)$ by the (nonzero) kernel average in (2.9), i.e.,

$$\hat{\rho}(y, x) = \frac{K(y, x)}{\hat{q}(y)}. \quad (2.10)$$

This type of normalization is called left normalization [10] due to the fact that the normalizer $\hat{q}(y)$ is evaluated at the “left” argument (y) of $K(y, x)$. In Section 3.3 ahead we will also consider alternative normalization strategies designed to control forecast biases. With either of these strategies, as $N \rightarrow \infty$ the function $\hat{\rho}(y, \cdot)$ converges for α -almost every starting point a_0 to a density

$$\rho(y, \cdot) = \frac{K(y, x)}{q(y)} \quad (2.11)$$

in $L^1(A, \alpha)$ such that $\int_A \rho(y, \pi(\cdot)) d\alpha = 1$. Moreover, $\rho(y, \cdot)$ is bounded by C/c .

Using $\hat{\rho}$ and ρ , we define the atomic probability measure

$$\hat{\nu}_y = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}(y, x_i) \delta_{a_i},$$

on (A, \mathcal{A}) and the measure ν_y such that

$$\nu_y(S) = \int_S \rho(y, \pi(\cdot)) d\alpha, \quad S \in \mathcal{A}.$$

Under the assumptions on K stated above, ν_y is absolutely continuous with respect to the invariant measure of the dynamics and $d\nu_y/d\alpha = \rho(y, \cdot) < \infty$. Moreover, it follows directly from the Birkhoff ergodic theorem that for every observable $f \in L^1(A, \alpha)$ the expectation value

$$\hat{F}(y) = \int_A f d\hat{\nu}_y = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}(y, x_i) f_i$$

converges to

$$F(y) = \int_A f d\nu_y = \int_A f \rho(y, \cdot) d\alpha$$

for α -almost every starting point a_0 .

In basic kernel analog forecasting, the forecast $\hat{F}_t(y)$ of f at lead time $t = k\tau$ given initial data y is obtained by evaluating the expectation of the time-shifted observable $U_t f$ from (2.4), i.e.,

$$\hat{F}_t(y) = \int_A U_t f d\hat{\nu}_y = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}(y, x_i) f_{i+k}. \quad (2.12)$$

In the limit of large data, $N \rightarrow \infty$, $\hat{F}_t(y)$ converges to

$$F_t(y) = \int_A U_t f d\nu_y. \quad (2.13)$$

Unlike traditional analog forecasting via (2.7), the kernel analog forecast in (2.12) depends continuously on the initial data due to the continuity of $K(\cdot, x)$. Moreover, the forecast value $\hat{F}_t(y)$ incorporates information from multiple samples in the training data. For this reason, we sometimes refer to kernel analog forecasts as ensemble analog forecasts, distinguishing them from single-analog approaches such as the conventional method in Section 2.2. As with (2.7), $\hat{F}_t(y)$ can be evaluated for arbitrary initial data in X . However, for y restricted to $M = \pi(A)$ there exists a function $\hat{f}_t \in L^1(A, \alpha)$ such that $\hat{f}_t(b) = \hat{F}_t(\pi(b))$. Conceptually, we think of \hat{f}_t as lying “above” \hat{F}_t . Similarly, there exists a function $f_t \in L^1(A, \alpha)$ lying above F_t .

For initial data of the form $y = \pi(b)$ with $b \in A$, the error in forecasts made via (2.12) is (cf. (2.8))

$$\hat{r}_t(b) = U_t f(b) - \hat{F}_t(y) = \int_A U_t f (d\delta_b - d\hat{\nu}_y),$$

and in the limit of large data that error becomes

$$r_t(b) = U_t f(b) - F_t(y) = \int_A U_t f (d\delta_b - d\nu_y).$$

Taking the expectation value of r_t with respect to the invariant measure of the dynamics, and using the fact that $\int_A U_t f d\alpha = \int_A f d\alpha = \bar{f}$, it follows that the expected forecast error over all initial conditions is given by

$$\bar{r}_t = \int_A r_t d\alpha = \bar{f} - \int_A \int_A \rho(y, x) U_t f(a) d\alpha(a) d\alpha(b), \quad (2.14)$$

where $x = \pi(a)$.

We now comment on the behavior of the expected error with respect to time, the mixing properties of the dynamics, and the properties of the kernel, assuming that f is a square-integrable observable in $L^2(A, \alpha)$.

First, note that \bar{r}_t is generally nonzero, even at $t = 0$ and for complete initial data. However, if the dynamical system is mixing, then \bar{r}_t approaches zero as $t \rightarrow \infty$ for both

complete and incomplete initial data. More specifically, mixing implies that for any two observables $f, g \in L^2(A, \alpha)$,

$$\lim_{t \rightarrow \infty} \int_A (U_t f) g \, d\alpha = \int_A f \, d\alpha \int_A g \, d\alpha,$$

and for any probability measure ν , absolutely continuous with respect to α ,

$$\lim_{t \rightarrow \infty} \int_A U_t f \, d\nu = \int_A f \, d\alpha.$$

Together, these two properties imply that under mixing dynamics the expected error vanishes at asymptotic times, $\lim_{t \rightarrow \infty} r_t = 0$. As expected, in this case, the mean squared error (MSE) $\|r_t\|^2$ approaches the ensemble variance of f , namely,

$$\lim_{t \rightarrow \infty} \|r_t\|^2 = \lim_{t \rightarrow \infty} \int_A r_t^2 \, d\alpha = \int_A (f - \bar{f})^2 \, d\alpha.$$

Analog forecasting via (2.12) becomes unbiased, $\bar{r}_t = 0$, at all times and irrespectively of the mixing properties of the dynamics and the completeness of the observations if $\rho(y, x)$ is bistochastic, that is

$$\int_M \rho(y, \cdot) \, d\mu = \int_M \rho(\cdot, x) \, d\mu = 1.$$

Indeed, in this case, the double integral in (2.14) is equal to \bar{f} , and r_t vanishes identically. While $\rho(y, x)$ constructed via (2.11) is generally not bistochastic, this observation motivates the use of density functions with the bistochastic property. In Section 3.3, we will discuss two approaches for constructing density functions which become bistochastic in the limit of large data, either via a careful kernel normalization [16], or a combination of normalization and kernel localization [17, 10].

Even with bistochastic density functions, however, basic kernel analog forecasting is overly diffusive and can suffer from large root mean squared error (RMSE). To see this, note that the forecast formula for $F_t(y) = f_t(b)$ can be expressed in terms of the action of an averaging operator $P : L^2(A, \alpha) \mapsto L^2(A, \alpha)$ such that

$$Pf(b) = \int_A p(b, a) f(a) \, d\alpha(a), \quad p(b, a) = \rho(\pi(b), \pi(a)).$$

In particular, we have $f_t(b) = PU_t f(b)$. Assume now that P is a “good” averaging operator for the dynamical system in the sense that it is positive-definite and self-adjoint on $L^2(A, \alpha)$, and α is an ergodic invariant measure of the Markov semigroup $\{P^n\}_{n \in \mathbb{N}}$. Assume also for simplicity that f has zero mean on $L^1(A, \alpha)$. Even in this favorable scenario, the RMSE $\|r_t\| = \|(I - P)U_t f\|$ cannot be less than $\gamma\|f\|$, where γ is the spectral gap of P . In other words, the fact that the forecast in (2.12) is determined by averaging $U_t f$ on A means that even if the initial data y are complete and happen to lie exactly on the training dataset, $F_t(y)$ will generally have nonzero error. The methods put forward in the following two sections attempt to address these deficiencies while maintaining the desirable features of basic kernel analog forecasting.

2.4. Kernel analog forecasting with Nyström extension

Our first refinement of basic kernel analog forecasting is based on the Nyström method for out-of-sample extension in reproducing kernel Hilbert spaces [18]. This method builds a data-driven, finite-dimensional Hilbert space of observables which can be evaluated exactly at arbitrary points on A using only information from the finite training dataset. If it happens that the prediction observable $U_t f$ actually lies in that Hilbert space for all lead times of interest, then Nyström extension essentially provides an error-free forecast for arbitrary initial data. Of course, the requirement that $U_t f$ lies entirely in this space will rarely hold in practice, and analog forecasting via the Nyström method will accrue errors and biases from the residual of $U_t f$ lying outside that finite-dimensional space. Nevertheless, the method is likely to perform well with data-driven observables (see Section 2.1) formed by linear combinations of leading kernel eigenfunctions, such as the patterns of low-frequency SST variability studied in Section 6.2.

2.4.1. Nyström extension for partially observed ergodic systems The starting point of analog forecasting with Nyström extension is a positive-semidefinite kernel $\sigma : X \times X \mapsto \mathbb{R}$. Similarly to the kernel K in Section 2.3, σ provides a measure of pairwise similarity between points in the observation space, but it is required to have somewhat different properties, as follows:

- (i) Symmetry: $\sigma(y, x) = \sigma(x, y)$ for all $x, y \in X$;
- (ii) Non-negativity: $\sum_{i,j=1}^n c_i c_j \sigma(x_i, x_j) \geq 0$ for all $n \in \mathbb{N}$, $x_1, \dots, x_n \in X$, and $c_1, \dots, c_n \in \mathbb{R}$;
- (iii) Boundedness: $\sigma \leq C$ for a number $C > 0$;

Property (iii), in conjunction with the fact that (M, \mathcal{M}, μ) is a probability space, implies that $\sigma(y, \cdot)$ and $\sigma(\cdot, x)$ are square-integrable functions in $L^2(M, \mu)$. Similarly, $\sigma(y, \pi(\cdot))$ and $\sigma(\pi(\cdot), x)$ are square-integrable functions in $L^2(A, \alpha)$. As with property (iii) of K in Section 2.3, the boundedness assumption on σ can be relaxed at the expense of increasing the complexity of the exposition below.

For our purposes, it is natural to work with positive-semidefinite kernels constructed from non-symmetric kernels K on $X \times X$ satisfying the conditions of Section 2.3. As with kernel normalization, there are several ways of carrying out this procedure, and as concrete examples we consider

$$\hat{\sigma}_L(y, x) = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}(x_i, y) \hat{\rho}(x_i, x), \quad \hat{\sigma}_R(y, x) = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}(y, x_i) \hat{\rho}(z, x_i),$$

where $\hat{\rho}(y, x_i)$ are the weights determined by left normalization of K via (2.10). We call the procedures to obtain $\hat{\sigma}_L$ and $\hat{\sigma}_R$ “left” and “right” symmetrizations, respectively. Properties (i) and (ii) follow directly by construction of these kernels, and property (iii) follows from the boundedness of $\hat{\rho}$ established in Section 2.3. In the limit of large data, $N \rightarrow \infty$, $\hat{\sigma}_L(y, x)$

and $\hat{\sigma}_R(y, x)$ respectively converge to the symmetric, positive semidefinite, bounded kernels

$$\sigma_L(y, x) = \int_M \rho(\cdot, y) \rho(\cdot, x) d\mu, \quad \sigma_R(y, x) = \int_M \rho(y, \cdot) \rho(x, \cdot) d\mu.$$

Note that, in general, the symmetric kernels introduced above are neither row- nor column-stochastic with respect to the invariant measure of the dynamics, i.e.,

$$\int_M \hat{\sigma}_L(y, \cdot) d\mu \neq 1, \quad \int_M \hat{\sigma}_L(\cdot, x) d\mu \neq 1,$$

and similarly for σ_L , $\hat{\sigma}_R$, and σ_R . In Section 3.3, we will discuss alternative normalization strategies of K [16, 10, 17] leading to bistochastic symmetric kernels in the limit of large data. Hereafter, when convenient we use the symbol $\hat{\sigma}$ to represent either of $\hat{\sigma}_L$ or $\hat{\sigma}_R$, and σ to represent either of σ_L or σ_R . Moreover, we use the shorthand notations

$$\hat{s}(b, a) = \hat{\sigma}(\pi(b), \pi(a)), \quad s(a, b) = \sigma(\pi(b), \pi(a)). \quad (2.15)$$

Because \hat{s} and s meet analogous conditions to (i)–(iii) above, they are symmetric positive-semidefinite kernels on $A \times A$. We write \hat{s}_L and s_L (\hat{s}_R and s_R) whenever we wish to distinguish between the kernels on $A \times A$ derived by left (right) symmetrization.

Having constructed a positive-semidefinite kernel with any of the above methods, our method follows closely the geometric harmonics technique of Coifman and Lafon [18], with the difference that we work with a partially observed system (i.e., our goal is to construct spaces of functions on A as opposed to $M = \pi(A)$), and we employ the kernel \hat{s} constructed from finite data (as opposed to s which is more closely related to the case studied in [18]). In particular, we turn attention to the reproducing kernel Hilbert space (RKHS) $\hat{\mathcal{H}}$ on M associated with \hat{s} . According to the Moore-Aronszajn theorem [1], this space is the unique Hilbert space of scalar-valued functions on M with inner product $\langle \cdot, \cdot \rangle_{\hat{\mathcal{H}}}$, such that (i) for every $b \in A$, the function $\hat{s}(b, \cdot)$ is in $\hat{\mathcal{H}}$; (ii) for every $f \in \hat{\mathcal{H}}$ and $b \in A$, point evaluation can be expressed in terms of the bounded linear operator

$$f \mapsto f(b) = \langle f, \hat{s}(b, \cdot) \rangle_{\hat{\mathcal{H}}}.$$

This last equation expresses the reproducing property of $\hat{\mathcal{H}}$, and it follows from this property that $\langle \hat{s}(b, \cdot), \hat{s}(a, \cdot) \rangle_{\hat{\mathcal{H}}} = \hat{s}(b, a)$. More generally, $\hat{\mathcal{H}}$ consists of functions of the form $f = \sum_{i=1}^{\infty} c_i \hat{s}(\cdot, b_i)$, with $\sum_{i,j=1}^{\infty} c_i c_j \hat{s}(b_i, b_j) < \infty$. The inner product between f and another function $f' = \sum_{i=1}^{\infty} c'_i \hat{s}(\cdot, b'_i)$ is $\langle f, f' \rangle_{\hat{\mathcal{H}}} = \sum_{i,j=1}^{\infty} c_i c'_j \hat{s}(b_i, b'_j)$. We define the RKHS \mathcal{H} associated with the kernel s and its inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ in an analogous manner, but note that this space cannot be constructed from finite datasets. Note that due to (2.15) all functions in $\hat{\mathcal{H}}$ and \mathcal{H} are constant on the pre-images $\pi^{-1}(x)$ of points in $x \in M$. Moreover, $f = \sum_{i=1}^{\infty} c_i \hat{s}(\cdot, b_i)$ lies above a unique function $F = \sum_{i=1}^{\infty} c_i \hat{\sigma}(\cdot, \pi(b_i))$ on X such that $f(b) = F(\pi(b))$ for any $b \in A$. The function F is in the RKHS on X associated with $\hat{\sigma}$.

Next, consider the Hilbert space $L^2(A, \hat{\alpha})$ equipped with the purely atomic measure $\hat{\alpha} = N^{-1} \sum_{i=0}^{N-1} \delta_{a_i}$ on the sampled points in M . This space consists of all measurable,

square-summable functions $f : M \mapsto \mathbb{R}$ such that $\int_A f^2 d\hat{\alpha} = N^{-1} \sum_{i=0}^{N-1} f^2(a_i) < \infty$, and we denote the inner product and norm of this space by

$$\langle f, g \rangle_{\hat{\alpha}} = \int_A fg d\hat{\alpha} = \frac{1}{N} \sum_{i=0}^{N-1} f(a_i)g(a_i), \quad \|f\|_{\hat{\alpha}} = \langle f, f \rangle_{\hat{\alpha}}^{1/2},$$

respectively. Since in an ergodic system the set of fixed points and periodic orbits has measure zero, the sampled states $\{a_0, \dots, a_N\}$ are all distinct with probability 1. Therefore, $L^2(A, \hat{\alpha})$ is naturally isomorphic to \mathbb{R}^N , equipped with the inner product $\langle f, g \rangle_{\mathbb{R}^N} = N^{-1} \sum_{i=0}^{N-1} f_i g_i$. Our analysis could equally be performed in this space, but we prefer to work with $L^2(A, \hat{\alpha})$ to emphasize that elements of $L^2(A, \hat{\alpha})$ and $L^2(A, \alpha)$ are functions on the same underlying space.

Because the set of singularities of any function in $L^2(A, \alpha)$ has α -measure zero, it follows that for α -almost any starting point a_0 , every f in $L^2(A, \alpha)$ is also in $L^2(A, \hat{\alpha})$. In particular, this is true for $\hat{s}(b, \cdot)$ for every $b \in A$. On the other hand, that f is in $L^2(A, \hat{\alpha})$ does not guarantee that it is also in $L^2(A, \alpha)$. Indeed, the set of singularities of $f \in L^2(A, \hat{\alpha})$ has $\hat{\alpha}$ measure zero, but can have nonzero α measure, and therefore $\int_A f^2 d\alpha$ can be infinite. If, however, $\|f\|_{\hat{\mathcal{H}}}$ is finite, it follows that $\|f\|$ and $\|f\|_{\hat{\alpha}}$ are both finite. Therefore, $\hat{\mathcal{H}}$ lies in the intersection of $L^2(A, \alpha)$ and $L^2(A, \hat{\alpha})$. Nyström extension techniques exploit this property to stably map subspaces of $L^2(A, \hat{\alpha})$ to subspaces of $L^2(A, \alpha)$. Due to the reproducing property, functions in these subspaces can be evaluated at arbitrary points on A without approximation.

To perform Nyström extension, we begin by introducing the integral operator $\hat{S} : L^2(A, \hat{\alpha}) \mapsto \hat{\mathcal{H}}$, mapping functions in the Hilbert space with the atomic sampling measure to the RKHS through the formula

$$\hat{S}f(b) = \int_A \hat{s}(b, \cdot) d\hat{\alpha} = \frac{1}{N} \sum_{i=0}^{N-1} \hat{s}(b, a_i) f(a_i).$$

The adjoint, $\hat{S}^* : \hat{\mathcal{H}} \mapsto L^2(A, \hat{\alpha})$, carries out the reverse operation, and it is a consequence of the reproducing property [18] that for any $g \in \hat{\mathcal{H}}$ and any $f \in L^2(A, \hat{\alpha})$,

$$\langle \hat{S}^*g, f \rangle_{\hat{\alpha}} = \langle g, \hat{S}f \rangle_{\hat{\mathcal{H}}} = \langle g, f \rangle_{\hat{\alpha}}.$$

This means that for $\hat{\alpha}$ -almost every a , $\hat{S}^*g(a) = g(a)$, or, equivalently, $\hat{S}^*g(a_i) = g(a_i)$ on the sampled states a_i . Due to this property, we interpret \hat{S}^* as a restriction operator from functions in $\hat{\mathcal{H}}$ to functions in $L^2(A, \hat{\alpha})$.

Consider now the positive-semidefinite self-adjoint operator $\hat{G} : L^2(A, \hat{\alpha}) \mapsto L^2(A, \hat{\alpha})$, $\hat{G} = \hat{S}^*\hat{S}$. By the spectral theorem, there exists an orthonormal basis $\{\hat{\phi}_0, \dots, \hat{\phi}_{N-1}\}$ of $L^2(A, \hat{\alpha})$ consisting of eigenfunctions of this operator. Denoting the corresponding, non-negative, eigenvalues by $\{\hat{\lambda}_0, \dots, \hat{\lambda}_{N-1}\}$, the eigenfunctions and eigenvalues satisfy the relation

$$\langle g, \hat{G}\hat{\phi}_k \rangle_{\hat{\alpha}} = \langle \hat{S}g, \hat{S}\hat{\phi}_k \rangle_{\hat{\mathcal{H}}} = \hat{\lambda}_k \langle g, \hat{\phi}_k \rangle_{\hat{\alpha}}$$

for any test function $g \in L^2(A, \hat{\alpha})$. That is, we have

$$\frac{1}{N} \sum_{i,j=0}^{N-1} g(a_i) \hat{s}(a_i, a_j) \hat{\phi}_k(a_j) = \hat{\lambda}_k \sum_{i=0}^{N-1} g(a_i) \hat{\phi}_k(a_i). \quad (2.16)$$

Writing

$$\hat{s}_L(a_i, a_j) = \frac{1}{N} \sum_{k=0}^{N-1} \hat{p}(a_k, a_i) \hat{p}(a_k, a_j), \quad \hat{s}_R(a_i, a_j) = \frac{1}{N} \sum_{k=0}^{N-1} \hat{p}(a_i, a_k) \hat{p}(a_j, a_k),$$

where $\hat{p}(a_i, a_j) = \hat{\rho}(\pi(a_i), \pi(a_j))$, the solutions of this eigenvalue problem can be found by computing the singular value decomposition of the $N \times N$ matrix $\mathbf{P} = [\hat{p}(a_i, a_j)/N]$. In the case of left symmetrization, we have $(\hat{\phi}_k(a_0), \dots, \hat{\phi}_k(a_{N-1}))^\top = \vec{v}_k$, where v_k is the k -th right singular vector of \mathbf{P} , and in the case of right symmetrization $(\hat{\phi}_k(a_0), \dots, \hat{\phi}_k(a_{N-1}))^\top = \vec{u}_k$, where \vec{u}_k is \mathbf{P} 's k -th left singular vector. In both cases, $\hat{\lambda}_k = \hat{\sigma}_k^2$, where $\hat{\sigma}_k$ is the k -th singular value of \mathbf{P} . In fact, in some cases, including $\hat{\rho}$ constructed via the methods of Section 3.3, $\mathbf{P} = \mathbf{D}^{-1}\mathbf{S}$ for a diagonal matrix \mathbf{D} and a symmetric matrix \mathbf{S} , and therefore it is a normal matrix. In such cases, the left and right singular vectors of \mathbf{P} coincide with its eigenvectors, and its singular values are equal to the absolute values of its eigenvalues.

Given a function $f = \sum_{i=0}^{N-1} c_i \hat{\phi}_i$ in $L^2(A, \hat{\alpha})$, it is a consequence of the reproducing property that $c_i = \langle \hat{\phi}_i, f \rangle_{\hat{\alpha}} = \hat{\lambda}_i \langle \hat{\phi}_i, f \rangle_{\hat{\mathcal{H}}}$. Moreover, given another function $g = \sum_{i=0}^{N-1} c'_i \hat{\phi}_i$ we have $\langle f, g \rangle_{\hat{\mathcal{H}}} = \sum_{i=0}^{N-1} c_i c'_i / \hat{\lambda}_i$. It therefore follows that the elements of $L^2(A, \hat{\alpha})$ which are in $\hat{\mathcal{H}}$ are all functions $f = \sum_{i=0}^{N-1} c_i \hat{\phi}_i$, $\sum_{i=0}^{N-1} c_i^2 < \infty$, such that $\|f\|_{\hat{\mathcal{H}}}^2 = \sum_{i=0}^{N-1} c_i^2 / \hat{\lambda}_i$ is finite. Intuitively, the kernel \hat{s} induces a measure of roughness of functions in $L^2(A, \hat{\alpha})$ through the corresponding RKHS norm; in particular, functions that do not project strongly on the eigenfunctions corresponding to small eigenvalues have small RKHS norm and are deemed smooth, or weakly oscillatory.

An important property of the eigenfunctions of \hat{G} with nonzero corresponding eigenvalues is that the associated functions

$$\hat{\psi}_i = \hat{\lambda}_i^{-1/2} \hat{S} \hat{\phi}_i, \quad \hat{\psi}_i \in \hat{\mathcal{H}}, \quad (2.17)$$

are orthonormal on $\hat{\mathcal{H}}$,

$$\langle \hat{\psi}_i, \hat{\psi}_j \rangle_{\hat{\mathcal{H}}} = (\hat{\lambda}_i \hat{\lambda}_j)^{-1/2} \langle \hat{S} \hat{\phi}_i, \hat{S} \hat{\phi}_j \rangle_{\hat{\mathcal{H}}} = \langle \hat{\phi}_i, \hat{\phi}_j \rangle_{\hat{\alpha}} = \delta_{ij},$$

and moreover these functions are uniquely defined at every point on A (as opposed to $\hat{\phi}_i$ whose values are arbitrary on sets of $\hat{\alpha}$ measure zero). It is straightforward to verify that $\hat{\psi}_i$ is an eigenfunction of the positive-semidefinite, self-adjoint operator $\hat{S} \hat{S}^* : \hat{\mathcal{H}} \mapsto \hat{\mathcal{H}}$ with corresponding eigenvalue $\hat{\lambda}_i$. Coifman and Lafon [18] call such functions geometric harmonics. Ordering the eigenvalues in decreasing order, $\hat{\lambda}_0 \geq \hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_{N-1}$, the set $\{\hat{\psi}_0, \dots, \hat{\psi}_{l-1}\}$ with $\hat{\lambda}_{l-1} > 0$ can be interpreted as the set of l orthonormal functions

in $\hat{\mathcal{H}}$ whose elements are maximally concentrated on the sampled points, in the sense of extremizing the Rayleigh quotient $\|f\|_{\hat{\mathcal{H}}}^2/\|f\|_{\hat{\alpha}}^2$.

While the $\hat{\psi}_i$ are guaranteed to be square-integrable functions by virtue of the fact that $\hat{\mathcal{H}} \subset L^2(A, \alpha)$, they are generally non-orthogonal on $L^2(A, \alpha)$. In particular, we have

$$\tilde{w}_{ij} = \langle \hat{\psi}_i, \hat{\psi}_j \rangle = \frac{(\hat{\lambda}_i \hat{\lambda}_j)^{-1/2}}{N^2} \sum_{k,l=0}^{N-1} \hat{\phi}_i(a_k) \tilde{s}_2(a_k, a_l) \hat{\phi}_j(a_l), \quad \tilde{s}_2(a_k, a_l) = \int_A \hat{s}(a_k, \cdot) \hat{s}(a_l, \cdot) d\alpha.$$

Therefore, $\langle \hat{\psi}_i, \hat{\psi}_j \rangle \neq 0$ for eigenfunctions corresponding to distinct eigenvalues unless the $(\hat{\phi}_i(a_0), \dots, \hat{\phi}_i(a_{N-1}))$ happen to be eigenvectors of the matrix $[\tilde{s}_2(a_k, a_l)]$. Note that \tilde{s}_2 , and therefore \tilde{w}_{ij} , cannot be computed exactly from finite data.

Despite their lack of orthogonality at finite N , the geometric harmonics converge to orthogonal functions on $L^2(A, \alpha)$ as $N \rightarrow \infty$, in the sense that their counterparts $\{\psi_0, \psi_1, \dots\}$ obtained from the kernel s are orthogonal on $L^2(A, \alpha)$. The procedure to obtain $\{\psi_i\}$ in the limit of large data is completely analogous to the finite N case, and can be carried out essentially by replacing all quantities with overhats in the above with their non-overhat counterparts. That is, instead of \hat{S} and \hat{G} , we have the kernel integral operators $S : L^2(A, \alpha) \mapsto \mathcal{H}$ and $G = S^*S : L^2(A, \alpha) \mapsto L^2(A, \alpha)$, respectively. In this case, G is a Hilbert-Schmidt operator, whose eigenfunctions $\{\phi_0, \phi_1, \dots\}$ provide an orthonormal basis of $L^2(A, \alpha)$, and the functions $f = \sum_{i=0}^{\infty} c_i \phi_i \in L^2(A, \alpha)$ which are also in \mathcal{H} satisfy $\|f\|_{\mathcal{H}}^2 = \sum_{i=0}^{\infty} c_i^2 / \lambda_i < \infty$. Moreover, the geometric harmonics $\psi_i = S\phi_i$ are orthonormal eigenfunctions of the operator $SS^* : \mathcal{H} \mapsto \mathcal{H}$, and it follows immediately that

$$w_{ij} = \langle \psi_i, \psi_j \rangle = (\lambda_i \lambda_j)^{-1/2} \langle S\phi_i, S\phi_j \rangle = (\lambda_i \lambda_j)^{-1/2} \langle \phi_i, G\phi_j \rangle_{\mathcal{H}} = \lambda_i \delta_{ij}. \quad (2.18)$$

In the limit of large data, $\tilde{w}_{ij} \rightarrow w_{ij}$ and the $\{\hat{\psi}_i\}$ become $L^2(A, \alpha)$ -orthogonal.

The lack of orthogonality of the geometric harmonics on $L^2(A, \alpha)$ can also be understood as a form of overfitting the training data. In particular, setting $\hat{w}_{ij} = \langle \hat{\psi}_i, \hat{\psi}_j \rangle_{\hat{\mathcal{H}}} = \hat{\lambda}_i \delta_{ij}$ and

$$\hat{s}_2(a_k, a_l) = \frac{1}{N} \sum_{j=0}^{N-1} \hat{s}(a_k, a_j) \hat{s}(a_l, a_j),$$

we write $\tilde{w}_{ij} = \hat{w}_{ij} + \Theta_{ij}$, where

$$\Theta_{ij} = \frac{(\hat{\lambda}_i \hat{\lambda}_j)^{-1/2}}{N^2} \sum_{k,l=0}^{N-1} \hat{\phi}_i(a_k) (\tilde{s}_2(a_k, a_l) - \hat{s}_2(a_k, a_l)) \hat{\phi}_j(a_l)$$

are quantities that vanish as $N \rightarrow \infty$. Thus, for a collection $\{\hat{\psi}_0, \dots, \hat{\psi}_{l-1}\}$ of geometric harmonics, we can study the matrices $\tilde{\mathbf{W}} = [\tilde{w}_{ij}]$, $\hat{\mathbf{W}} = [\hat{w}_{ij}]$, and $\mathbf{\Theta} = [\Theta_{ij}]$. Here, the ratio of the matrix norms $\|\mathbf{\Theta}\|/\|\hat{\mathbf{W}}\|$ measures how strongly the inner product relationships of $\{\hat{\psi}_0, \dots, \hat{\psi}_{l-1}\}$ on the Hilbert space $L^2(A, \hat{\alpha})$ associated with the sampling measure differ from those on the Hilbert space $L^2(A, \alpha)$ associated with the invariant measure. Large

values of this ratio signify that $\{\hat{\psi}_0, \dots, \hat{\psi}_{l-1}\}$ overfits the training data. Another way of characterizing overfitting is through the ratio $\kappa(\tilde{\mathbf{W}})/\kappa(\hat{\mathbf{W}})$ of the condition numbers of $\tilde{\mathbf{W}}$ and $\hat{\mathbf{W}}$. Qualitatively, for fixed N , we expect both $\|\Theta\|/\|\hat{\mathbf{W}}\|$ and $\kappa(\tilde{\mathbf{W}})/\kappa(\hat{\mathbf{W}})$ to increase with l , so there is a tradeoff between the explanatory power afforded by using large l and the increased risk of overfitting the training data. We will return to this point in Section 2.4.2 ahead.

2.4.2. Analog forecasting We now have the necessary ingredients to perform analog forecasting with Nyström extension. First, setting a bandwidth parameter l , we construct an l -dimensional subspace of $L^2(A, \alpha)$ spanned by the first l geometric harmonics, $\hat{B}_l = \text{span}\{\hat{\psi}_0, \dots, \hat{\psi}_{l-1}\}$. Intuitively, we think of \hat{B}_l as a space of bandlimited observables with respect to the kernel \hat{s} . Three key properties of an observable $\eta = \sum_{i=0}^{l-1} c_i \hat{\psi}_i \in \hat{B}_l$ are:

- (i) The expansion coefficients c_i can be determined without approximation from the values of η on the finite training dataset by evaluating

$$c_i = \langle \eta, \hat{\psi}_i \rangle_{\hat{\mathcal{H}}} = \lambda_i^{-1/2} \langle \eta, \hat{\phi}_i \rangle_{\hat{\alpha}} = \frac{\lambda_i^{-1/2}}{N} \sum_{j=0}^{N-1} \eta(a_j) \hat{\phi}_i(a_j); \quad (2.19)$$

- (ii) η can be evaluated at arbitrary points on A using

$$\eta(b) = \sum_{i=0}^{l-1} c_i \hat{\psi}_i(b) = \frac{1}{N} \sum_{i=0}^{l-1} \sum_{j=0}^{N-1} \hat{s}(b, a_j) c_i \hat{\lambda}_i^{-1/2} \hat{\phi}_i(a_j);$$

- (iii) η lies above a unique function $H : X \mapsto \mathbb{R}$ given by

$$H(y) = \frac{1}{N} \sum_{i=0}^{l-1} \sum_{j=0}^{N-1} \hat{\sigma}(y, x_j) c_i \hat{\lambda}_i^{-1/2} \hat{\phi}_i(a_j),$$

and we have $\eta(b) = H(\pi(b))$.

Note that, restricted to $M \subseteq X$, H is a square-function in $L^2(M, \mu)$. In what follows, we denote the function on X lying above $\hat{\psi}_i$ by $\hat{\Psi}_i$.

Next, consider the prediction observable $f \in L^2(A, \alpha)$. Fixing a time interval $[0, t_1]$, we compute the expansion coefficients

$$c_i(t) = \langle \hat{\psi}_i, U_t f \rangle_{\hat{\mathcal{H}}} = \frac{1}{N} \sum_{j=0}^{N-1} U_t f(a_j) \hat{\phi}_i(a_j) = \frac{1}{N} \sum_{j=0}^{N-1} f_{i+k} \hat{\phi}_i(a_j) \quad (2.20)$$

for all $i \in \{0, \dots, l-1\}$ and $t = k\tau \in [0, t_1]$. Then, given an initial condition $y \in X$, we define the Nyström analog forecast as

$$\hat{F}_t(y) = \sum_{i=0}^{l-1} c_i(t) \hat{\Psi}_i(y). \quad (2.21)$$

For initial data of the form $y = \pi(b)$, $b \in A$, the forecast error with this scheme is

$$\hat{r}_t(b) = \hat{F}_t(y) - U_t f(b).$$

To interpret this forecast, suppose first that $U_t f$ lies in \hat{B}_l for all discrete times $t = k\tau \in [0, t_1]$. Then, we have $U_t f = \sum_{i=0}^{l-1} c_i(t) \hat{\psi}_i$, and for initial data of the form $y = \pi(b)$ the forecast function $\hat{F}_t(y)$ equals $U_t f(b)$ and the forecast error vanishes. Note that there is no approximation in this result despite the fact that the training dataset contains a finite number of samples.

Of course, the assumption that $U_t f$ is a bandlimited observable in \hat{B}_l for all lead times of interest is unlikely to hold in practice. First, in many forecast scenarios the prediction observable is pre-determined, and the chances of it lying entirely in a finite-dimensional data-driven subspace are small. In the language of Section 2.1, this situation would likely arise for objectively defined observables. Even if $U_t f$ were in \hat{B}_l at a fixed time (say at $t = 0$), \hat{B}_l is generally not an invariant subspace of the dynamics, and there is no reason to expect $U_t f$ to lie in this space at other times.

Given a general prediction observable $f \in L^2(A, \alpha)$, then with probability 1 (i.e., excluding cases of α -measure zero where f is singular on the sampled states a_i) we can write

$$U_t f = \eta_{t,l} + \hat{r}_{t,l},$$

where $\eta_{t,l} = \sum_{i=0}^{l-1} c_i(t) \hat{\psi}_i$ is in \hat{B}_l , and $\hat{r}_{t,l} \in B_l^\perp$ is a residual in the orthogonal complement of \hat{B}_l in $L^2(A, \hat{\alpha})$. The coefficients $c_i(t)$ can be computed without approximation using (2.20), and moreover the refinement $\eta_{t,l+1} - \eta_{t,l}$ at each l is $L^2(A, \hat{\alpha})$ -orthogonal to the previous approximation η_l . Thus,

$$\|\hat{r}_{t,l+1}\|_{\hat{\alpha}}^2 = \|\hat{r}_{t,l}\|_{\hat{\alpha}}^2 - \|\eta_{t,l+1} - \eta_{t,l}\|_{\hat{\alpha}}^2, \quad (2.22)$$

and the residual norm on $L^2(A, \hat{\alpha})$ is a non-increasing function of l , i.e., $\|\hat{r}_{t,l+1}\|_{\hat{\alpha}} \leq \|\hat{r}_{t,l}\|_{\hat{\alpha}}$.

However, this does not mean that the residual norm $\|\hat{r}_{t,l+1}\|$ on $L^2(A, \alpha)$, which is equal to the RMSE of the forecast for initial data drawn from the invariant measure, is also a non-increasing function of l . Indeed, since the $\hat{\psi}_i$ are non-orthogonal on $L^2(A, \alpha)$ (see Section 2.4.1), the refinement at each l is not $L^2(A, \alpha)$ -orthogonal to the previous approximation, and instead of (2.22) we have

$$\|\hat{r}_{t,l+1}\|^2 = \|\hat{r}_{t,l}\|^2 - \|\eta_{t,l+1} - \eta_{t,l}\|^2 + 2\langle \eta_l, \eta_{l+1} - \eta_l \rangle, \quad (2.23)$$

which is not guaranteed to be non-increasing. An increasing residual norm with l is interpreted as overfitting the training data. According to (2.18), in the limit of large data the geometric harmonics become $L^2(A, \alpha)$ -orthogonal, and the expected residual becomes a non-increasing function of l .

Equation (2.23) can also be expressed in terms of the Gramm matrix elements $\tilde{w}_{ij} = \hat{w}_{ij} + \Theta_{ij}$ introduced in Section 2.4.1, namely,

$$\|\hat{r}_{t,l+1}\|^2 = \|\hat{r}_{t,l}\|^2 - \hat{\lambda}_{l+1} c_{l+1}^2 + 2 \sum_{i=1}^l c_i c_{l+1} \Theta_{i,l+1}.$$

From the above, we see that the possibility of increasing expected residual with l is related to the non-orthogonality of the geometric harmonics with respect to the invariant measure of the dynamics. At fixed sample number N , the coefficients Θ_{ij} generally increase in magnitude with i and j , and the risk of increase of $\|r_{t,l}\|$ also increases. In other words, we are faced with the usual tradeoff between the ability to predict large classes of observables (large dimension of \hat{B}_l) and risk of overfitting the training data. In the applications of Section 6, we will select values of l achieving a good balance between these competing sources of error by cross-validation against a dataset independent of the training data. A more detailed study of such strategies is beyond the scope of the present paper, but should be an interesting avenue of future research.

Despite the potential shortcomings of the Nyström method when dealing with objectively-defined observables, in several interesting applications the objective is to actually construct data-driven observables that capture the salient features of high-dimensional data and have favorable predictability properties. Eigenfunctions of kernels designed specifically for analyzing dynamical system data are theoretically [6, 30] and experimentally [32, 31, 14, 56] known to perform well in terms of timescale separation and physical interpretability. Therefore, observables lying in the bandlimited spaces associated with these kernels are natural candidates for prediction. An attractive property observables in these spaces is that they have vanishing reconstruction error at forecast initialization (i.e., $\hat{F}_0(y) = f(b)$ for all $y = \pi(b)$). In Section 3.2, we will argue that these spaces may persist (i.e., be approximately invariant) under the dynamical flow, improving the predictability of the corresponding data-driven observables on short to intermediate times.

2.5. Kernel analog forecasting with Laplacian pyramids

In this Section, we present a kernel analog forecasting technique based on the Laplacian pyramids algorithm for multiscale approximation of functions introduced by Rabin and Coifman [50] and further improved by Fernández et al. [28]. This forecasting approach combines aspects of the basic kernel analog forecasting technique in Section 2.3 with RKHS ideas from Section 2.4. In what follows, we begin by presenting the approach in the limit of large data (Section 2.5.1), and then discuss its implementation for finite datasets (Section 2.5.2). Throughout this Section, we assume that the prediction observable f and all kernel-derived densities are square-integrable functions in $L^2(A, \alpha)$.

2.5.1. Laplacian pyramids in the limit of large data First, consider the basic kernel analog forecast $F_t(y)$ in (2.13), determined by the expectation of $U_t f$ with respect to the kernel-dependent probability measure ν_y . For initial data of the form $y = \pi(b)$, $b \in A$, this expression can alternatively be expressed in terms of an averaging operator, $P_0 : L^2(A, \alpha) \mapsto L^2(A, \alpha)$, namely

$$f_{t,0}(b) = P_0 U_t f(b) = \int_A p_0(b, a) U_t f(a) d\alpha(a). \quad (2.24)$$

Here, $p_0(b, a) = \rho_0(y, \pi(a))$, and ρ_0 is a density determined by normalization of a kernel K on $X \times X$, e.g., via (2.11) so that $\int_A p_0(b, \cdot) d\alpha = 1$ and $\int_M \rho_0(y, \cdot) d\mu = 1$ for all $b \in A$ and $y \in X$. (The reason for introducing the “0” subscripts will become apparent below.) Moreover, as discussed in Section 2.3, the error in $f_{t,0}$ is given by

$$r_{t,0}(b) = U_t f(b) - P_0 U_t f(b).$$

For an observable with mean $\bar{f} = \int_A f d\alpha = \int_A U_t f d\alpha$ we can also write

$$f_{t,0}(b) = \bar{f} + P_0 U_t f'(b), \quad r_{t,0}(b) = U_t f'(b) - P_0 U_t f'(b),$$

where $f' = f - \bar{f}$.

While the Laplacian pyramids procedure can be operationally carried out for arbitrary averaging operators, in order to place bounds on the RMS forecast error, and ensure stability of the algorithm, we require that $\|(I - P_0)f\| \leq \|f\|$. This condition can be met for arbitrary observables in $L^2(A, \alpha)$ if P_0 is positive-semidefinite and self-adjoint on $L^2(A, \alpha)$, and α is an ergodic invariant measure of the Markov semigroup $\{P_0^n\}_{n \in \mathbb{N}}$ (see also Section 2.3). That is, we must have $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n P_0^k f = \int_A f d\alpha$ for all $f \in L^2(A, \alpha)$. The ergodicity of P_0 is equivalent to the requirement that $P_0 f = f$ implies that f is α -almost everywhere constant. Moreover, that P_0 is self-adjoint on $L^2(A, \alpha)$ implies that the kernel p_0 is symmetric and therefore bistochastic, $\int_A p_0(b, \cdot) d\alpha = \int_A p_0(\cdot, a) d\alpha = 1$. Such kernels can be constructed from more general similarity kernels via the normalization procedure developed in [16], which we will also employ in Section 3.3.3.

These assumptions lead to the residual bound

$$\|r_{t,0}\| = \|(I - P_0)U_t f'\| \leq \|f'\|, \quad (2.25)$$

where we have used the fact that $\|U_t f'\| = \|f'\|$. Moreover, the expected residual $\int_A r_{t,0} d\alpha$ vanishes due to the bistochasticity of p_0 (see Section 2.3), and the residual norm $\|r_{t,0}\|$ is equal to the RMS forecast error. Note if $U_t f'$ is not in the nullspace of P_0 the inequality in (2.25) becomes strict, i.e., $\|r_{t,0}\| < \|f'\|$. Furthermore, because $p_0(b, \cdot) = \rho_0(\pi(b), \pi(\cdot))$, the function $f_{t,0} : A \mapsto \mathbb{R}$ lies above a function $F_{t,0}$ on X given by

$$F_{t,0}(y) = \int_A \rho_0(y, \pi(a)) U_t f(a) d\alpha(a).$$

This function agrees with $f_{t,0}(b)$ for initial data of the form $y = \pi(b)$, and its restriction on $M = \pi(A)$ is a square-integrable function in $L^2(M, \mu)$.

While basic analog forecasting would stop at (2.24), Laplacian pyramids treats $r_{t,0}$ as another function which is to be approximated by averaging. In particular, noticing that $P_0 U_t f$ lies in a RKHS $\mathcal{H}_0 \subset L^2(A, \alpha)$ associated with the kernel p_0 , it adjusts the forecast by adding a correction, $g_{t,1}$, which lies in a RKHS \mathcal{H}_1 associated with a kernel p_1 .

There are several ways of constructing \mathcal{H}_1 , and more generally a sequence $\{\mathcal{H}_0, \mathcal{H}_1, \dots, \mathcal{H}_l\}$ of RKHSs such that $\mathcal{H}_i \subset L^2(A, \alpha)$. In the setting of exponentially

decaying, bandwidth-dependent kernels used in this work, one can simply fix the functional form of the kernel K_ϵ (e.g, use the radial Gaussian family in (2.3)), and set the \mathcal{H}_i to the RKHSs associated with the bistochastic kernels p_i constructed from K_{ϵ_i} through the procedure in Section 3.3.3 for a decreasing sequence $\epsilon_0 > \epsilon_1 > \dots > \epsilon_l$. Refs. [50, 28] use successive subdivisions of the form $\epsilon_i = \epsilon_0/2^i$, which we also employ here. Due to the bistochasticity of the p_i , all of the corresponding averaging operators $P_i : L^2(A, \alpha) \mapsto L^2(A, \alpha)$ are self-adjoint, and the range of P_i is contained in \mathcal{H}_i .

Having constructed a nested-RKHS family, we begin by approximating the residual at iteration 1 by acting on it with P_1 :

$$g_{t,1}(b) = P_1 r_{t,0}(b) = \int_A p_1(b, a) r_{t,1}(a) d\alpha(\alpha).$$

The function $g_{t,1}$ is now in \mathcal{H}_1 , and the residual $r_{t,1}(b) = r_{t,0} - g_{t,1}$ can be bounded in the $L^2(A, \alpha)$ norm through

$$\|r_{t,1}\| \leq \|r_{t,0}\| \leq \|f'\|.$$

We then define the first-level approximation of $U_t f$ as $f_{t,1} = f_{t,0} + g_{t,1}$, and this function is in \mathcal{H}_1 . Since both p_0 and p_1 are bistochastic, this approximation has vanishing expected error, $\int_A r_{t,1} d\alpha = 0$, and $\|r_{t,1}\|$ is equal to the RMS forecast error.

Iterating this procedure l times yields the l -level Laplacian pyramids approximation of the forecast observable f at lead time t :

$$f_{t,l} = f_{t,0} + \sum_{i=1}^l g_{t,i}, \quad f_{t,l} \in \mathcal{H}_l, \quad (2.26)$$

where $f_{t,0}$ is given by (2.24), and $g_{t,i}$ are functions in \mathcal{H}_i determined iteratively through the formulas

$$g_{t,i} = P_i r_{t,i-1}, \quad r_{t,i} = \begin{cases} f_t - f_{t,0}, & i = 0, \\ r_{t,i-1} - g_{t,i}, & i \geq 1. \end{cases} \quad (2.27)$$

This approximation has the residual norm

$$\|r_{t,l}\| \leq \|r_{t,l-1}\| \leq \dots \leq \|r_{t,0}\| \leq \|f'\|, \quad (2.28)$$

and the above inequalities become strict if the $r_{t,i}$ and $U_t f'$ are not in the nullspaces of the corresponding P_i . Moreover, $\int_A r_{t,i} d\alpha$ vanishes by bistochasticity of the p_i , and $\|r_{t,i}\|$ is equal to the RMS forecast error. As with $f_{t,0}$, $f_{t,l}$ lies above a function $F_{t,l} : X \mapsto \mathbb{R}$ such that $F_{t,l}(\pi(b)) = f_{t,l}(b)$, and the restriction of $F_{t,l}$ on M is in $L^2(M, \mu)$.

In general, the limiting behavior of the algorithm as $l \rightarrow \infty$ depends on whether or not the limit space in the sequence $\{\mathcal{H}_l\}$ is the whole $L^2(A, \alpha)$ space. Since the kernels associated with \mathcal{H}_l have the form $K_{\epsilon_l}(\pi(\cdot), \pi(\cdot))$, functions in \mathcal{H}_l are always constant on the pre-images $\pi^{-1}(x)$ of points $x \in M$. Therefore, if π is not invertible on its image (i.e., the observations are incomplete), \mathcal{H}_l will converge to a strict subspace of $L^2(A, \alpha)$ and $\lim_{l \rightarrow \infty} \|r_{t,l}\|$ will generally be nonzero.

Note that while the requirement that the P_i are self-adjoint on the L^2 space associated with the invariant measure of the dynamics is very natural, the method can still be applied for more general averaging operators with ergodic invariant measures α_i . In such cases, the ranges of the P_i are contained in RKHSs with reproducing kernels given by left or right symmetrization of p_i (see Section 2.4.1). Moreover, residual bounds analogous to (2.28) can be derived for the $L^2(A, \alpha_i)$ norms. If the α_i are sufficiently smooth (e.g., absolutely continuous) relative to the invariant measure of the dynamics, then it may be possible to obtain from these bounds useful bounds for the residual norm on $L^2(A, \alpha)$. Here, we do not study the behavior of the residual in more detail, but we note the result in [28] that for complete observations, radial Gaussian kernels, and uniform sampling density, the residual norm decreases with l faster than any algebraic rate.

2.5.2. Laplacian pyramids for finite datasets Analog forecasting with Laplacian pyramids and finite datasets is structurally similar to the formulation in the limit of large data, but one has to additionally confront the risk of overfitting. More specifically, in the finite-sample case we replace the probability kernels p_i (which involve integrals with respect to the invariant measure; see (3.15)) with bistochastic kernels $\hat{p}_i = \hat{\rho}_i(\pi(\cdot), \pi(\cdot))$ on $L^1(A, \hat{\alpha})$ constructed using the method in Section 3.3.3. These kernels satisfy $\hat{p}_i(b, a) = \hat{p}_i(a, b)$ for arbitrary points in A , and moreover $N^{-1} \sum_{j=0}^{N-1} \hat{p}_i(b, a_j) = 1$. While \hat{p}_i is not bistochastic on $L^1(A, \alpha)$, it converges to the bistochastic kernel p_i from Section 2.5.1 as $N \rightarrow \infty$.

Using the same sequence of decreasing bandwidths ϵ_i as in Section 2.5.1, we construct the sequence of RKHSs $\{\hat{\mathcal{H}}_0, \hat{\mathcal{H}}_1, \dots, \hat{\mathcal{H}}_l\}$, $\hat{\mathcal{H}}_i \subset L^2(A, \hat{\alpha})$, associated with \hat{p}_i , and introduce the positive-semidefinite, self-adjoint averaging operators $\hat{P}_i : L^2(A, \hat{\alpha}) \mapsto L^2(A, \hat{\alpha})$ such that

$$\hat{P}_i f(b) = \int_A \hat{p}_i(b, a) f(a) d\hat{\alpha}(a) = \frac{1}{N} \sum_{j=0}^{N-1} \hat{p}_i(b, a_j) f(a_j), \quad \hat{P}_i f \in \hat{\mathcal{H}}_i.$$

Note that by square-integrability of \hat{p}_i , the $\hat{\mathcal{H}}_i$ are also subspaces of $L^2(A, \alpha)$ (see Section 2.4.1). In what follows, we will assume that $\hat{\alpha}$ is an ergodic invariant measure of the Markov semigroups $\{\hat{P}_i^n\}_{n \in \mathbb{N}}$; this condition will be fulfilled so long as the $N \times N$ Markov matrices $\hat{\mathbf{P}}_i = [\hat{p}_i(a_j, a_k)/N]$ on the training data are irreducible.

At the zeroth level of the Laplacian pyramids construction, we evaluate

$$\hat{f}_{t,0} = \hat{P}_0 U_t f.$$

This approximation has a residual in $L^2(A, \hat{\alpha})$ given by

$$\hat{r}_{t,0} = U_t f - \hat{P}_0 U_t f,$$

which we approximate in turn through

$$\hat{g}_{t,1} = \hat{P}_1 \hat{r}_{t,0}.$$

From this point onward, the iterations are repeated l times analogously to (2.26) and (2.27), replacing P_i by \hat{P}_i , until the Laplacian pyramids approximation

$$\hat{f}_{t,l} = \hat{f}_{t,0} + \sum_{i=1}^l \hat{g}_{t,i}$$

is completed. The function $\hat{f}_{t,l}$ lies above a unique function $\hat{F}_{t,l}$ on X constructed analogously to $F_{t,l}$ in Section 2.5.1.

By self-adjointness of the \hat{P}_i on $L^2(A, \hat{\alpha})$ and ergodicity of the associated diffusion semigroups, the residual norm $\|\hat{r}_{t,l}\|_{\hat{\alpha}}$ on $L^2(A, \hat{\alpha})$ is a non-increasing function of l . In particular, we have

$$\|\hat{r}_{t,l}\|_{\hat{\alpha}} \leq \|\hat{r}_{t,l-1}\|_{\hat{\alpha}} \leq \dots \leq \|\hat{r}_{t,0}\|_{\hat{\alpha}} \leq \|U_t f - \hat{f}_t\|_{\hat{\alpha}}, \quad (2.29)$$

where $\hat{f}_t = \int_A U_t f d\alpha$ the expectation of the time-shifted prediction observable with respect to the sampling measure. Note that unlike the corresponding expectation value with respect to the invariant measure (used, e.g., in (2.28)), \hat{f}_t is not time-independent (but converges to a time-independent function in the limit of large data). As with (2.28), the inequalities in (2.29) become strict if the $\hat{r}_{t,i}$ and $U_t f$ are not in the nullspaces of the corresponding \hat{P}_i operators.

Of course, that the residual norm on $L^2(A, \hat{\alpha})$ is non-increasing does not imply that the residual norm on $L^2(A, \alpha)$ is non-increasing too. In fact, at fixed N , we generally expect the discrepancy between $\|\hat{r}_{t,l}\|_{\hat{\alpha}}$ and $\|\hat{r}_{t,l}\|$ to increase with l . Intuitively, as l increases $\hat{\mathcal{H}}_l$ contains functions of progressively smaller scales, which are prone to similar overfitting effects as those discussed in the context of the Nyström extension in Section 2.4.

In the case of Laplacian pyramids, the occurrence of overfitting is made particularly transparent from the fact that the probability measures induced on (A, \mathcal{A}) by the \hat{p}_i converge as $i \rightarrow \infty$ to pullbacks of δ -measures in the data space. To examine this in more detail, fix a state $a_j \in A$ in the training dataset and consider the probability measure $\hat{\nu}_{a_j,i}$ on (A, \mathcal{A}) such that $\hat{\nu}_{a_j,i}(S) = \int_S \hat{p}_i(a_j, \cdot) d\hat{\alpha}$. For typical bandwidth-dependent kernels (including the radial Gaussian kernel in (2.3) and the kernels of Section 3), as $i \rightarrow \infty$ and $\epsilon_i \rightarrow 0$, this probability measure converges to $\Delta_{a_j} = \delta_{\pi(a_j)}$. That is, as $i \rightarrow \infty$, \hat{P}_i acts on f by averaging over its values on the points in A that are mapped to a single point $x_j = \pi(a_j)$ in M . Moreover, if the observations are complete, a_j is the only point in A which is sent to x_j and $\Delta_{a_j} = \delta_{a_j}$.

To alleviate this issue, Fernández et al. [28] advocate replacing \hat{p}_i by the “hollow” transition probability kernels \tilde{p}_i such that

$$\tilde{p}_i(b, a_j) = \tilde{\rho}_i(y, x_j), \quad \tilde{\rho}_i(y, x_j) = \begin{cases} 0, & y = x_j, \\ \hat{\rho}_i(y, x_j)/\zeta(y), & y \neq x_j. \end{cases} \quad (2.30)$$

In the above, $y = \pi(b)$, $x_j = \pi(a_j)$, and $\zeta(y) = \sum_{x_k \neq y} \hat{\rho}_i(y, x_k)/\tilde{N}(y)$, where $\tilde{N}(y)$ is equal to the number of points in the training dataset $\{x_k\}$ different from y . Inspired by leave-one-out

cross-validation techniques, this modification of the kernel has a small effect at large ϵ_i , but prevents the probability measures $\hat{\nu}_{a_j,i}$ to degenerate to Δ_{a_j} as $\epsilon_i \rightarrow 0$ (at the expense of introducing a small bias in the invariant measure of the corresponding Markov semigroup). As a result, the residual norm $\|\hat{r}_{t,l}\|_{\hat{\alpha}}$ may increase with l . In the experiments of Section 6, we will always use the transition kernels \tilde{p}_i when making kernel analog forecasts with Laplacian pyramids. Operationally, we find that as l grows, $\|r_{t,l}\|_{\hat{\alpha}}$ initially decays, but eventually reaches a minimum and then starts increasing again. Choosing l as the minimizer of $\|r_{t,l}\|_{\hat{\alpha}}$ appears to be an effective practical criterion for balancing increased explanatory power with risk of overfitting the training data.

3. Choice of kernel

While the kernel analog forecasting techniques presented in Section 2 place few constraints on the form of the kernel, the choice of kernel will ultimately have a strong impact on forecast skill. In this Section we discuss specific examples of kernels for analog forecasting which we will subsequently employ in the applications of Section 6. Here, we focus on two aspects of kernel design, namely (i) the selection of the functional form of the kernel K in Section 2.3, and (ii) the normalization and symmetrization of that kernel to construct the various out-of-sample extension and averaging operators in Sections 2.4 and 2.5.

In this discussion, we restrict attention to so-called local kernels [10], i.e., kernels with an exponential decay controlled by a bandwidth parameter, ϵ , such as the radial Gaussian kernel K_ϵ in (2.3). Intuitively, ϵ controls the “resolution” of K_ϵ in the phase space A (e.g., through the associated nested RKHSs of Section 2.5), and one can study the asymptotic behavior of these kernels as $\epsilon \rightarrow 0$. In several cases of interest (e.g., when X is a Hilbert space and M is a smooth manifold embedded in X), the limiting behavior of local kernels is governed by well-studied operators on manifolds such as diffusion operators and vector fields. Thus, one can use the properties of these operators to inform kernel design.

3.1. Delay-coordinate maps

Delay-coordinate maps [57, 13, 52] is a powerful tool for recovering topological features of attractors of dynamical systems lost through partial observations. Here, we use delay-coordinate maps as a means of actually constructing the data space X on which subsequent kernel calculations are performed.

Consider the same dynamical system $(A, \mathcal{A}, \Phi_t, \alpha)$ as Section 2.1, but observed through a different observation map, $\tilde{\pi} : A \mapsto Z$, where Z is a data space. We assume that Z and $\tilde{\pi}$ have similar measurability properties as those required to hold for X and π in Section 2.1. Fixing a positive integer parameter q (the number of delays), we map the point a in phase space A to a point x in $X = Z^q$ (hereafter, called delay-coordinate space) via one of the two

mappings

$$\begin{aligned} a \mapsto x = \pi_-(a) &= (\tilde{\pi}(a), \tilde{\pi}(\Phi_\tau^{-1}(a)), \dots, \tilde{\pi}(\Phi_{(q-1)\tau}^{-1}(a))), \\ a \mapsto x = \pi_+(a) &= (\tilde{\pi}(a), \tilde{\pi}(\Phi_\tau(a)), \dots, \tilde{\pi}(\Phi_{(q-1)\tau}(a))), \end{aligned}$$

Note that in the non-invertible (semigroup) case π_- is not defined, and this will have an impact on the utility of delay-coordinate mapping in analog forecasting of these systems. In what follows, whenever we use π_- we will tacitly assume that the dynamics is invertible. Denoting the time series of observations in Z by $\{z_0, z_1, \dots\}$, where $z_i = \tilde{\pi}(a_i)$, the action of π_- and π_+ on a state a_i sampled in the training phase is determined by the forward or backward lagged sequences,

$$x_i = \pi_-(a_i) = (z_i, z_{i-1}, \dots, z_{i-(q-1)}), \quad x_i = \pi_+(a_i) = (z_i, z_{i+1}, \dots, z_{i+(q-1)}), \quad (3.1)$$

respectively. Notice that so long as the time ordering of the training data is available $\pi_-(a_i)$ and $\pi_+(a_i)$ can be evaluated without explicit knowledge of the underlying state a_i .

According to a theorem of Takens [57], which was generalized in [52] to systems with fractal attractors under prevalent sets of observation functions, if Φ_t is a flow on a compact attractor A of a dynamical system, then with high probability the points $x_i = \pi(a_i)$ lie on a set $M = \pi(A) \subseteq X$ which is in one-to-one correspondence with A (i.e., A and M are diffeomorphic manifolds), provided that q is sufficiently large. Thus, because knowledge of the position on the attractor is sufficient to determine the future evolution of the system, the time series $\{x_i\}$ becomes Markovian even if $\{z_i\}$ is non-Markovian.

Since analog forecasting relies on following the evolution of observables on such time series starting from given initial conditions (analogs), it is natural to identify those initial conditions using the lagged sequences x_i rather than z_i . Even in situations where M and A are not in one-to-one correspondence despite delay-coordinate mapping, we expect π to have recovered at least some predictive information which is not contained in $\tilde{\pi}$. For this reason, we will construct our kernels for analog forecasting in delay-coordinate space X , as opposed to the “snapshot” space Z . This approach should be beneficial in both traditional analog forecasting with respect to Euclidean distances and the kernel-based methods developed here. Note that the special case $q = 1$ corresponds to no embedding, $X = Z$.

Before proceeding, it is important to note a key difference between the backward (π_-) and forward (π_+) delay-coordinate maps. Namely, π_- can be evaluated in a “real-time” fashion by concatenating previously observed snapshots, whereas π_+ cannot be evaluated without using information from the future. As a result, in analog forecasting scenarios where the snapshots z_i used in forecast initialization are obtained incrementally, the smallest lead time for which the time interval needed to evaluate π_+ does not overlap with the forecast time interval is $t_0 = (q-1)\tau$. In other words, when measuring the skill of forecasts made with π_+ , the first $q-1$ leads are “used up” for forecast initialization. Due to these considerations, when the dynamics are invertible π_- is preferable for analog forecasting (whereas π_- and π_+ perform equally well for topological recovery). When the dynamics is not invertible and π_-

cannot be used there is a tradeoff between topological recovery (favored by large q) and ease of forecast initialization (favored by small q).

3.2. Local kernels with dynamics-adapted features

3.2.1. Local kernels in delay-coordinate space Hereafter, we assume that the data space Z is a Hilbert space with inner product $\langle z, z' \rangle_Z$ and norm $\|z\| = \langle z, z \rangle_Z^{1/2}$. Under this assumption, the delay-coordinate space X naturally acquires a Hilbert space structure with inner product given by summing the $\langle \cdot, \cdot \rangle_Z$ inner products at each lag. That is, for $x = (\zeta_0, \dots, \zeta_{q-1})$ and $z' = (\zeta'_0, \dots, \zeta'_{q-1})$ with $\zeta_i, \zeta'_i \in Z$ we define $\langle x, x' \rangle = \sum_{i=0}^{q-1} \langle \zeta_i, \zeta'_i \rangle_Z$ and $\|x\| = \langle x, x \rangle^{1/2}$. An example of a kernel in delay-coordinate space is the radial Gaussian kernel

$$K_\epsilon(y, x) = e^{-D^2(y, x)/\epsilon}, \quad D(y, x) = \|y - x\|, \quad \epsilon > 0. \quad (3.2)$$

This kernel is equivalent to (2.3) with the distance function $D(y, x)$ set to the norm $\|y - x\|$ of delay-coordinate space.

As with other kernels encountered in Section 2, K_ϵ induces an integral operator, $\mathcal{K}_\epsilon : L^2(A, \alpha) \mapsto L^2(A, \alpha)$ acting on square-integrable functions on A through the formula

$$\mathcal{K}_\epsilon f(b) = \int_A K_\epsilon(\pi(b), \pi(a)) f(a) d\alpha(a). \quad (3.3)$$

An important property of kernels defined in delay-coordinate space such as (3.2) is that their associated integral operators depend on the dynamical system generating the data. In particular, because each point in delay-coordinate space corresponds to a segment of dynamical evolution, different dynamical systems will induce different kernel values, and hence different operators, even when observed in the same observation space. We refer to kernels whose properties depend on the dynamical system generating the data as “dynamics-adapted” kernels.

In the case that $M = \pi(A) \subseteq X$ is a smooth, compact manifold diffeomorphic to A , K_ϵ is symmetric, and ϵ is small, the influence of the dynamics on K_ϵ and \mathcal{K}_ϵ can be studied from the point of view of a Riemannian geometry induced on A from the ambient geometry of X . Here, the intuition is that as for a fixed reference point b , as ϵ becomes small $K_\epsilon(\pi(b), \pi(\cdot))$ is only appreciable in an exponential neighborhood of b (see, e.g., (3.10) ahead). In this neighborhood, the behavior of \mathcal{K}_ϵ is governed by the leading integral moments of $K_\epsilon(\pi(b), \pi(\cdot))$ on A [10], which are governed in turn by the leading derivatives of the distance function. In particular, the second derivative (Hessian) of $D(\pi(b), \pi(\cdot))$ is closely related to the Riemannian metric tensor inherited by M from the ambient space inner product.

In practical applications, we are interested in the behavior of integral operators defined for the atomic sampling measure $\hat{\alpha}$. That is, instead of (3.2) we consider, $\hat{\mathcal{K}}_\epsilon : L^2(A, \hat{\alpha}) \mapsto L^2(A, \hat{\alpha})$, where

$$\hat{\mathcal{K}}_\epsilon f(b) = \int_A K_\epsilon(\pi(b), \pi(a)) f(a) d\hat{\alpha}(a). \quad (3.4)$$

For the same reasons as those leading to overfitting in Laplacian pyramids (see Section 2.5), as ϵ decreases at fixed N , the properties of this operator become increasingly different from those of \mathcal{K}_ϵ . Moreover, the smaller ϵ is, the larger N must be to achieve good consistency between $\hat{\mathcal{K}}_\epsilon$ and \mathcal{K}_ϵ [53]. Intuitively, at small ϵ , \mathcal{K}_ϵ is sensitive to small-scale features of f , requiring a large number of samples to approximate the integral in (3.3) by (3.4). In the limit $\epsilon \rightarrow 0$, infinitely many samples are required. Thus, in practice, the limit $\epsilon \rightarrow 0$ is realized as a limit of large data. As mentioned in Section 2.3, additional complications can arise if M is noncompact. In such cases, a variable-bandwidth variant of (3.2) can be used to ensure stable behavior as $\epsilon \rightarrow 0$ and $N \rightarrow \infty$ [8].

In [6], Berry et al. established that under the assumption that Takens' theorem holds (i.e., A and M are diffeomorphic manifolds), increasing the number of delays q biases the kernel-induced Riemannian metric of A towards an invariant subspace associated with the most stable Lyapunov direction of the dynamical system (specifically, the most stable Oseledets subspace associated with the Lyapunov metric [49]). As has been demonstrated in a number of applications [32, 6, 14, 56], the geometry of the data in delay-coordinate space significantly enhances the ability to extract intrinsic dynamical timescales with kernel eigenfunctions. Thus, it may be desirable to perform delay-coordinate mapping even if the observations are full to take advantage of this property. Below, we will argue that the Riemannian geometry of data in delay-coordinate space is also beneficial for analog forecasting.

3.2.2. Cone kernels Besides (2.3), several other kernels for dynamical systems have been proposed in the literature [54, 32, 6, 58, 30, 8] and could be employed in the analog forecasting techniques developed here in different scenarios. In this paper, we work with the family of “cone kernels” introduced in [30]. A novel aspect of these kernels (which will turn out to be beneficial for analog forecasting) is that they feature an explicit dependence on the time tendency of the data, estimated through finite differences in time. For instance, $v = (x_i - x_{i-1})/\tau$ corresponds to a first-order backward finite-difference approximation for the time derivative of the data, but higher-order and/or central or forward schemes can equally be used. In [30], a geometrical interpretation was given for these quantities in terms of the vector field $V = d\Phi_t/dt|_{t=0}$ on the attractor generating the dynamics (assuming again that A and M are diffeomorphic manifolds). As a result, incorporating v in kernels allows one to bias the Riemannian metric of the data in a way that depends explicitly on the generator.

In cone kernels, that dependence enters through the norm of V , estimated through the norm $\|v\|$ in ambient data space, and the angle between V and the relative displacement vector $\omega = y_i - x_j$ of time-indexed samples in delay coordinate space. The cosine of that angle is estimated through $\cos \eta = \langle v, \omega \rangle / (\|v\| \|\omega\|)$, and similarly we compute $w = (y_i - y_{i-1})/\tau$ and $\cos \theta = \langle w, \omega \rangle / (\|w\| \|\omega\|)$. Introducing a parameter $\zeta \in [0, 1)$ (in addition to the kernel

bandwidth ϵ), we define the cone kernel by

$$K_{\epsilon, \zeta}(y_i, x_j) = e^{-L_\zeta(y_i, x_j)/\epsilon}, \quad L_\zeta(y_i, x_j) = \frac{\|\omega\|^2}{\|w\|\|v\|} [(1 - \zeta \cos^2 \theta)(1 - \zeta \cos^2 \eta)]^{1/2}. \quad (3.5)$$

Note that as with radial the Gaussian kernels in (3.2), cone kernels are symmetric, but the function L_ζ appearing in the argument of the exponential is anisotropic when $\zeta > 0$.

As ζ approaches 1, this kernel assigns increasingly higher affinity to data samples whose relative displacement vector is aligned with either v and/or w ; i.e, it assigns strong affinity to pairs of samples whose relative displacement vector lies within a narrow cone aligned with the dynamical vector field. This leads to distance contraction along the dynamical flow, such that the new averaging operators will be well-adapted to extract observables varying on intrinsic slow timescales. In particular, it can be shown [30] that in the induced geometry of cone kernels the norm $\|u\|_\zeta$ of tangent vectors on the manifold is related to the ambient-space norm $\|u\|$ through the expression

$$\|u\|_\zeta^2 = \frac{1}{\|V\|^2} \left(\|u\|^2 - \zeta \frac{\langle V, u \rangle^2}{\|V\|^2} \right). \quad (3.6)$$

(We emphasize again that that this results requires that the A and $M = \pi(A)$ are diffeomorphic manifolds; thus, in the case that $\tilde{\pi}$ is an incomplete observation map, we also require that Takens' theorem holds.) For $\zeta > 0$ the metric preferentially contracts tangent vectors aligned with V having $\langle V, u \rangle^2 / \|V\|^2 \approx \|u\|^2$, and the amount of length contraction becomes arbitrarily large as $\zeta \rightarrow 1$. As a result, in the cone-kernel geometry with $\zeta \approx 1$, small spherical balls centered on the initial data y will preferentially contain samples in the training dataset lying along a tube of integral curves of the dynamical vector field containing the integral curve passing through y . Integral operators constructed from cone kernels (e.g., (3.3) and the averaging operators of Section 3.3 ahead) will therefore have weaker impact on slow observables varying weakly on these neighborhoods aligned with the dynamical flow. On the other hand, fast observables exhibiting strong variability along the orbits of the dynamics will be strongly averaged. Thus, integral operators constructed from cone kernels can be thought of as intrinsic low-pass filters associated with the dynamical system.

Another important property of the cone kernel construction is that it is stable under changes in observation modality, in the sense that the associated integral operators derived from datasets related by invertible transformations will be similar even if the corresponding ambient-space metrics (i.e., the dot product and corresponding norm in (3.6)) are significantly different. We also note that as a result of the scaling by $1/(\|w\|\|v\|)$ cone kernels acquire a scale invariance which is useful for computing affinities for heterogeneous datasets consisting of measurements with different physical dimension (units) [14]. The outcome of this scaling factor in the induced geometry is a conformal scaling by the inverse squared norm of the dynamical vector field which can be seen in (3.6). Heuristically, the result of this conformal factor is to preferentially increase the resolution (i.e., to “zoom in”) at the

regions of the data manifold where the phase space velocity $\|V\|$ is small), and decrease the resolution when $\|V\|$ is large. Large $\|V\|$ (and small sampling density) is expected to occur during metastable regime transitions which will be a prominent feature of the atmospheric model studied in Section 6.1. Therefore, the conformal transformation by $1/(\|w\|\|v\|)$ should improve the stability of kernel integral operators used to analyze such datasets. In [32], the atmospheric model of Section 6.1 was studied using the kernel in (3.5) with $\zeta = 0$, and it was found that this transformation is highly beneficial for recovering the intermittent regime transitions in Galerkin-reduced models.

3.3. Kernel normalization

As discussed in Section 2, the raw kernel K_ϵ must be normalized prior to use in analog forecasting. In Section 2.3, we described a simple normalization strategy that converts $K_\epsilon(y, \cdot)$ to a density function $\hat{\rho}_\epsilon(y, \cdot)$ on the training dataset through left normalization. Here, we outline two alternative normalization strategies which produce bistochastic kernels in the limit of large data, hence eliminating forecast biases. Throughout this Section, we restrict attention to symmetric kernels with $K_\epsilon(y, x) = K_\epsilon(x, y)$. Most of the results of this Section are valid for general symmetric local kernels, and for notational simplicity we drop the ζ subscript from cone kernels unless we are making reference to a result specific to this kernel family.

As stated in the preamble to Section 3, we are interested the behavior of these normalization strategies and the associated averaging operators acting on functions on A in the limit of large data and small kernel bandwidth. In the case of averaging operators constructed through symmetric kernels, and for complete observations (i.e., A is diffeomorphic to its image $M = \pi(A)$ under the observation map), the behavior of the averaging operators can be characterized through second-order self-adjoint differential operators generating a diffusion process on A . In the incomplete observations case, the limiting behavior of the averaging operators is significantly harder to analyze, but nevertheless they remain useful for unbiased analog forecasting.

3.3.1. Diffusion maps normalization The first normalization strategy used in this work was introduced in the diffusion maps algorithm of Coifman and Lafon [17], and further studied by Berry and Sauer in [10]. Here, instead of left-normalizing K_ϵ directly (as done in (2.10)), we first perform a right normalization,

$$\hat{K}'_\epsilon(y, x) = \frac{K_\epsilon(y, x)}{\hat{q}_\epsilon^{1/2}(x)}, \quad (3.7)$$

and then left-normalize \hat{K}'_ϵ to produce a row-stochastic kernel:

$$\hat{\rho}_\epsilon(y, x) = \frac{\hat{K}'_\epsilon(y, x)}{\hat{d}_\epsilon(y)}, \quad \hat{d}_\epsilon(y) = \frac{1}{N} \sum_{i=0}^{N-1} \hat{K}'_\epsilon(y, x_i). \quad (3.8)$$

For completeness, we note that in the original formulation of diffusion maps [17], the right normalization in (3.7) is performed using arbitrary powers of \hat{q}_ϵ . Here, we have elected to normalize K_ϵ by $\hat{q}_\epsilon^{1/2}$ as this choice leads to orthogonal eigenfunctions with respect to the invariant measure of the dynamics.

By the pointwise ergodic theorem, as $N \rightarrow \infty$, the quantities \hat{q}_ϵ , \hat{K}'_ϵ , \hat{d}_ϵ , and $\hat{\rho}_\epsilon$ converge to

$$q_\epsilon(y) = \int_M K_\epsilon(y, \cdot) d\mu, \quad K'_\epsilon(y, x) = \frac{K_\epsilon(y, x)}{q_\epsilon^{1/2}(y)}, \quad d_\epsilon(y) = \int_M K'_\epsilon(y, \cdot) d\mu, \quad \rho_\epsilon(y, x) = \frac{K'_\epsilon(y, x)}{d_\epsilon(y)}.$$

Moreover, $\hat{\rho}_\epsilon(y, \cdot)$ and $\rho_\epsilon(y, \cdot)$ are densities on $L^2(M, \hat{\mu})$ and $L^2(M, \mu)$, respectively, so that

$$\int_M \hat{\rho}_\epsilon(y, \cdot) d\hat{\mu} = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}_\epsilon(y, x_i) = 1, \quad \int_M \rho_\epsilon(y, \cdot) d\mu = 1.$$

Setting $\hat{p}_\epsilon(b, a) = \hat{\rho}_\epsilon(\pi(b), \pi(a))$ and $p_\epsilon(b, a) = \rho_\epsilon(b, a)$, we also have

$$\int_A \hat{p}_\epsilon(b, \cdot) d\hat{\alpha} = \frac{1}{N} \sum_{i=0}^{N-1} \hat{p}_\epsilon(b, a_i) = 1, \quad \int_A p_\epsilon(b, \cdot) d\alpha = 1, \quad (3.9)$$

and the techniques of Sections 2.3–2.5 can be implemented using these densities.

Here, we are interested in the behavior of $\hat{\rho}_\epsilon$ and ρ_ϵ and the associated averaging operators $\hat{P}_\epsilon : L^2(A, \hat{\alpha}) \mapsto L^2(A, \hat{\alpha})$ and $P_\epsilon : L^2(A, \alpha) \mapsto L^2(A, \alpha)$, given by

$$\hat{P}_\epsilon f(b) = \int_A \hat{p}_\epsilon(b, a) f(a) d\hat{\alpha}(a), \quad P_\epsilon f(b) = \int_A p_\epsilon(b, a) f(a) d\alpha(a).$$

First, note that in general P_ϵ is not self-adjoint on $L^2(A, \alpha)$. However, expressing $P_\epsilon f(b)$ as

$$P_\epsilon f(b) = \int_A s_\epsilon(b, a) f(a) \frac{d_\epsilon(x)}{q_\epsilon(x)} d\alpha(a), \quad s_\epsilon(b, a) = \frac{K_\epsilon(b, a)}{d_\epsilon(\pi(b)) d_\epsilon(\pi(a))},$$

where $y = \pi(b)$, and $x = \pi(a)$, it is evident that P_ϵ is positive-semidefinite, and self-adjoint on $L^2(A, \alpha_\epsilon)$ for the measure α_ϵ with $d\alpha_\epsilon/d\alpha = d_\epsilon/q_\epsilon$. By the properties of K_ϵ listed in Section 2.3, α_ϵ is absolutely continuous with respect to α . Similarly, \hat{P}_ϵ is not self-adjoint on $L^2(A, \hat{\alpha})$, but it can be expressed as

$$\hat{P}_\epsilon f(b) = \int_A \hat{s}_\epsilon(b, a) f(a) \frac{\hat{d}_\epsilon(x)}{\hat{q}_\epsilon(x)} d\hat{\alpha}(a), \quad \hat{s}_\epsilon(b, a) = \frac{K_\epsilon(b, a)}{\hat{d}_\epsilon(\pi(b)) \hat{d}_\epsilon(\pi(a))}.$$

Thus, \hat{P}_ϵ is self-adjoint on $L^2(A, \hat{\alpha}_\epsilon)$ for the atomic measure $\hat{\alpha}_\epsilon$ with density $\hat{d}_\epsilon/\hat{q}_\epsilon$ relative to $\hat{\alpha}$.

Consider now the behavior of P_ϵ as $\epsilon \rightarrow 0$ assuming that $M = \pi(A)$ is diffeomorphic to A . Under this condition, due to the exponential decay of the kernel it is possible to approximate kernel integrals using

$$\frac{1}{\epsilon^{m/2}} \int_A K_\epsilon(\pi(b), \pi(a)) f(a) d\alpha(a) = \gamma(b) f(b) + O(\epsilon), \quad (3.10)$$

where m is the dimension of A and $\gamma(b)$ is a smooth positive function that does not depend on f or ϵ (e.g., [17, 10, 30]). (More specifically, $\gamma(b)$ depends on the density of α relative to the volume form of the Riemannian metric induced by the kernel.) Note that the completeness of the observations is essential to the validity of this approximation. In particular, if π is not invertible on its image, the kernel will fail to localize to an m -dimensional ball of arbitrarily small radius as $\epsilon \rightarrow 0$, invalidating (3.10).

Using (3.10) and the fact that $K_\epsilon(y, x) = K_\epsilon(x, y)$, we find

$$\int_A p_\epsilon(b, a) d\alpha(b) = 1 + O(\epsilon).$$

This means that with the diffusion maps normalization, $p_\epsilon(b, a)$ becomes asymptotically bistochastic as $\epsilon \rightarrow 0$. Moreover, $d\alpha_\epsilon/d\alpha \rightarrow 1$, and P_ϵ becomes self-adjoint on $L^2(A, \alpha)$. It is straightforward to verify that $p_\epsilon(b, a)$ is also bistochastic in the limit $\epsilon \rightarrow \infty$ where K_ϵ attains the constant value 1. By the pointwise ergodic theorem, $\lim_{N \rightarrow \infty} \hat{P}_\epsilon f(a) = P_\epsilon f(a)$ for α -almost every a and starting state a_0 is the training data, so in the limit of large data \hat{P}_ϵ acquires the self-adjointness properties of P_ϵ on $L^2(A, \alpha)$.

It follows from these results that, with the diffusion maps normalization, the basic kernel analog forecast through (2.13) becomes unbiased in the limit of large data and for $\epsilon \rightarrow 0$ and $\epsilon \rightarrow \infty$, in the sense that the expected forecast error in (2.14) over initial conditions drawn from the invariant measure on A vanishes. (In fact, that would also be the case for the simpler kernel normalization via (2.11).) However, at intermediate values of ϵ , $\int_A p_\epsilon(b, a) d\alpha(a) \neq 1$ and p_ϵ is not bistochastic (though it is always row-stochastic in accordance with (3.9)), and the expected error can be nonzero even as $N \rightarrow \infty$.

Of course, the fact that basic kernel analog forecasting with large-bandwidth kernels becomes unbiased does not have high practical significance, since such forecasts are likely to be highly uninformative to begin with. That property, however, turns out to be important in the case of analog forecasting with Laplacian pyramids where one sweeps through a broad range of bandwidths. Here, we have seen that the diffusion maps normalization improves the forecast biases via Laplacian pyramids over the basic left normalization at small ϵ and when the observations are complete (i.e., M and A are diffeomorphic manifolds and (3.10) holds), but biases are likely to remain at intermediate ϵ values (and at both intermediate and small ϵ values if the observations are incomplete). In Section 3.3.3, we will discuss an alternative normalization strategy [16] leading to unbiased forecast as $N \rightarrow \infty$ for arbitrary ϵ and observation map π .

3.3.2. Diffusion eigenfunctions and data-driven observables We now examine the impact of the diffusion maps normalization to the eigenfunctions of the averaging operators P_ϵ and \hat{P}_ϵ and the Nyström-extension based forecasting scheme in Section 2.4.

First, assuming that the observations are complete, it is possible to show [17, 10] that, uniformly on A ,

$$P_\epsilon f(b) = f(b) - \epsilon \mathcal{L}f(b) + O(\epsilon^2), \quad (3.11)$$

where $\mathcal{L} : L^2(A, \alpha) \mapsto L^2(A, \alpha)$ is the (positive semidefinite) generator of a stochastic gradient flow on A for a Riemannian metric g (that depends on the kernel) and a potential determined by the density θ of the invariant measure of the dynamics relative to the volume form of g . Specifically, we have

$$\mathcal{L}f = \Delta f - \frac{\Delta\theta^{1/2}}{\theta^{1/2}}f, \quad (3.12)$$

where $\Delta f = -\operatorname{div}_g \operatorname{grad}_g f$ is the Laplace-Beltrami operator associated with g .

Equation (3.11) establishes the pointwise convergence of $\epsilon^{-1}(I - P_\epsilon)f$ to $\mathcal{L}f$ as $\epsilon \rightarrow 0$. By ergodicity we can conclude that $\hat{\mathcal{L}}f(a) = \lim_{\epsilon \rightarrow 0} \epsilon^{-1} \lim_{N \rightarrow \infty} (I - \hat{P}_\epsilon)f(a)$ is α -almost surely equal to $\mathcal{L}f(a)$ too. However, spectral convergence, i.e., convergence of the eigenvalues and eigenfunctions of $\epsilon^{-1}(I - P_\epsilon)$ and $\epsilon^{-1}(I - \hat{P}_\epsilon)$ to those of \mathcal{L} is a more subtle question which, to our knowledge, has not been addressed in the literature for the diffusion maps normalization used here to construct P_ϵ and \hat{P}_ϵ . However, a number of related results suggest that spectral convergence may hold for the class of averaging operators employed in this work. In particular, in [4] Belkin and Niyogi establish spectral convergence of the graph Laplacian (an operator analogous to $\hat{\mathcal{L}}$ in our notation) to the Laplace-Beltrami operator for radial Gaussian kernels and uniform sampling density relative to the Riemannian measure. Moreover, in [63] von Luxburg et. al establish spectral convergence of normalized kernel integral operators which have structural similarities with the operators P_ϵ and \hat{P}_ϵ studied here. On the basis of these results, it appears plausible that spectral convergence holds for the operators employed in this work too. In what follows, we provide a geometrical interpretation of the eigenfunctions of P_ϵ and \hat{P}_ϵ , and the associated geometric harmonics, assuming spectral convergence of $\epsilon^{-1}(I - P_\epsilon)$ and $\epsilon^{-1}(I - \hat{P}_\epsilon)$ to \mathcal{L} .

In particular, an important property of the eigenfunctions of \mathcal{L} is that they are the extrema of the ratio

$$r(f) = \frac{E_g(f)}{\|f\|^2}, \quad E_g(f) = \int_A \|\operatorname{grad}_g f\|^2 d\alpha, \quad (3.13)$$

where $E_g(f)$ is a weighted Dirichlet energy measuring the expected roughness of f with respect to the Riemannian metric g and the invariant measure of the dynamics. Denoting the eigenvalues and eigenfunctions of \mathcal{L} by (κ_i, φ_i) , we have $0 = \kappa_0 < \kappa_1 \leq \kappa_2 \leq \dots \uparrow \infty$, $\langle \varphi_i, \varphi_j \rangle = \delta_{ij}$, and $\kappa_i = E_g(\varphi_i)$. Moreover, $P_\epsilon \varphi_i = \lambda_{\epsilon,i} \varphi_i + O(\epsilon^2)$, where $\lambda_{\epsilon,i} = e^{-\epsilon \kappa_i}$. It therefore follows that as $\epsilon \rightarrow 0$, the leading l eigenfunctions of P_ϵ have a geometrical interpretation as the least rough orthonormal set of l functions in $L^2(A, \alpha)$ in the sense of (3.13). By ergodicity, the eigenfunctions of \hat{P}_ϵ acquire the same property in the limits $\epsilon \rightarrow 0$ and $N \rightarrow \infty$. Below, we will see that for the choice of kernels in Section 3.2 this geometrical interpretation also has a dynamical interpretation which motivates the use of the φ_i as data-driven prediction observables.

A similar geometrical interpretation can be made for the geometric harmonics employed in kernel analog forecasting with Nyström extension. In particular, recall that in

the geometric harmonics construction as formulated in Section 2.4.1 one computes the eigenvalues and eigenfunctions, $(\lambda_{\epsilon,i}, \phi_{\epsilon,i})$, of a symmetric kernel integral operator $G_\epsilon : L^2(A, \alpha) \mapsto L^2(A, \alpha)$ constructed from left or right symmetrization of p_ϵ . Noticing that in the case of left symmetrization $G_{L,\epsilon} = P_\epsilon^* P_\epsilon$ and in the case of right symmetrization $G_{R,\epsilon} = P_\epsilon P_\epsilon^*$, where the adjoints are taken on $L^2(A, \alpha)$, the eigenfunctions of $G_{L,\epsilon}$ ($G_{R,\epsilon}$) are given by the right (left) singular vectors of P_ϵ , and the eigenvalues are given by the squares of the corresponding singular values. As $\epsilon \rightarrow 0$, P_ϵ becomes self-adjoint on $L^2(A, \alpha)$, and the $\phi_{\epsilon,i}$ converge to φ_i for both left and right symmetrization. Moreover, the eigenvalues become $\lambda_{\epsilon,i} = e^{-2\epsilon\kappa_i} + O(\epsilon^2)$. Thus, as ϵ becomes small, the geometric harmonic $\psi_{\epsilon,i} = S_{\epsilon,i}\phi_{\epsilon,i}$ corresponding to ϕ_i is approximated by $\psi_{\epsilon,i} \approx e^{-\epsilon\kappa_i}\varphi_i$. It therefore follows that for small ϵ (and thus large N), the spaces of bandlimited observables \hat{B}_l used in Section 2.4.2 for analog forecasting have a geometrical interpretation as the l -dimensional subspaces of $L^2(A, \alpha)$ consisting of the least rough observables in the sense of (3.13). Note also that for two functions $f = \sum_{i=0}^\infty c_i \phi_{\epsilon,i}$ and $f' = \sum_{i=0}^\infty c'_i \phi_{\epsilon,i}$ in the RKHS \mathcal{H}_ϵ associated with p_ϵ we have $\langle f, f' \rangle_{\mathcal{H}_\epsilon} \approx \sum_{i=0}^\infty c_i c'_i / e^{-2\epsilon\kappa_i}$.

Consider now the geometrical properties of the eigenfunctions of P_ϵ for the dynamics-adapted kernels of Section 3.2. First, for kernels constructed in delay-coordinate space, including (3.2) and (3.5), as more delays are added the data-driven observables constructed from the eigenfunctions are expected to become increasingly biased towards stable Lyapunov directions of the system [6]. That is, given an orthonormal basis $\{e_1, \dots, e_m\}$ of the tangent space at $a \in A$ with respect to a Lyapunov metric on A , ordered in order of increasing Lyapunov exponent, we heuristically expect that $f \in B_l = \text{span}\{\psi_0, \dots, \psi_{l-1}\}$ will vary predominantly along e_1, \dots, e_{m_s} , where m_s is the dimension of the most stable Oseledets subspace. That is, we expect that $\|\text{grad}_s f\| \gg \|\text{grad}_{s^\perp}(f)\|$ where grad_s and grad_{s^\perp} denote the orthogonal projection of the gradient along the subspaces spanned by $\{e_1, \dots, e_{m_s}\}$ and $\{e_{m_s+1}, \dots, e_m\}$, respectively. By invariance of the Oseledets subspaces, we also expect this property to hold for a reasonably long time under dynamical evolution. As a result, prediction observables initially lying in B_l are expected to remain in this space at least over moderately long times, alleviating the risk of bias in analog forecasting with Nyström extension (see Section 2.4.2).

Next, consider the eigenfunctions from the cone kernels in (3.5). In this case, it is possible to show that a consequence of the length contraction along the dynamical flow as ζ approaches 1, the corresponding Laplace-Beltrami operator admits the asymptotic expansion [30],

$$\Delta f = -\frac{\text{div}_g(V(f)V)}{1-\zeta} + O((1-\zeta)^0), \quad (3.14)$$

where in the above equation $V(f)$ is the directional derivative of f along the dynamical vector field. Note that a general Laplace-Beltrami operator $\Delta f = -\text{div grad } f$ (approximated, e.g., through a radial Gaussian kernel) depends on the full gradient of f , which generally depends on the observation map for the dynamical system through the ambient-space induced metric. On the other hand, the directional derivative $V(f)$ is intrinsic to the dynamical system, and

is independent of the observation map π . It also follows from (3.14) that as $\zeta \rightarrow 1$ the Dirichlet energy in (3.13) reduces to $E_g(f) = (1 - \zeta)^{-1} \int_A (V(f))^2 d\alpha + O((1 - \zeta)^0)$. This functional assigns large roughness to functions with large directional derivative $V(f)$, i.e., functions that exhibit strong variability along the dynamical flow. Equivalently, observables which are strongly bandlimited with respect to cone kernels are expected to vary slowly in the course of dynamical evolution. Such observables are good candidates for revealing intrinsic slow dynamical timescales, and are also expected to have favorable predictability properties.

3.3.3. Bistochastic kernels The diffusion maps normalization in Section 3.3.1 eliminates analog forecast biases with averaging operators at small and large bandwidths, but does not address biases at intermediate bandwidths. Here, we present an alternative normalization technique that eliminates biases at arbitrary bandwidths (in fact, the method is applicable for general kernels without bandwidth parameters). Our approach is based on the bistochastic kernel construction of Coifman and Hirn [16], which we adapt here to the case of finite datasets consisting of time-ordered observations of ergodic dynamical systems.

Recall that the diffusion maps normalization consists of the right normalization of the raw kernel in (3.7), followed by a left normalization in (3.8). Here, we carry out this procedure in the opposite sense. Specifically, we first left-normalize K_ϵ to construct the density function $\hat{\rho}_\epsilon(y, x)$ via (2.10), and then make a right normalization as follows:

$$\hat{\beta}_\epsilon(y, x) = \frac{\hat{\rho}_\epsilon(y, x)}{\hat{\omega}_\epsilon^{1/2}(x)}, \quad \hat{\omega}_\epsilon(x) = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}_\epsilon(x_i, x).$$

In the limit of large data, these quantities converge to

$$\beta_\epsilon(y, x) = \frac{\rho_\epsilon(y, x)}{\omega_\epsilon^{1/2}(x)}, \quad \omega_\epsilon(x) = \int_M \rho(y, x) d\mu(x).$$

To construct bistochastic kernels from $\hat{\beta}_\epsilon$ and β_ϵ , we symmetrize these kernels using right symmetrization as in Section 2.4.1. That is, we set

$$\hat{\sigma}_\epsilon(y, x) = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\beta}_\epsilon(y, x_i) \hat{\beta}_\epsilon(x, x_i), \quad \sigma_\epsilon(y, x) = \int_M \beta_\epsilon(y, x') \beta_\epsilon(x, x') d\mu(x'), \quad (3.15)$$

and define the corresponding symmetric kernels on $A \times A$ through $\hat{s}_\epsilon(b, a) = \hat{\sigma}_\epsilon(\pi(b), \pi(a))$ and $s_\epsilon(b, a) = \sigma_\epsilon(\pi(b), \pi(a))$. Clearly, \hat{s}_ϵ and s_ϵ are bistochastic on $L^1(A, \alpha)$ and $L^1(A, \hat{\alpha})$ since $\hat{\beta}_\epsilon$ and β_ϵ are column stochastic. Of course \hat{s}_ϵ is neither row- nor column-stochastic on $L^1(A, \alpha)$ but it converges to a bistochastic kernel in the limit of large data. By virtue of the fact that this result holds for arbitrary bandwidths, the averaging operators \hat{P}_ϵ constructed from \hat{s}_ϵ yield asymptotically unbiased analog forecasts when used in conjunction with the Laplacian pyramids algorithm for both complete and incomplete observations.

4. Error metrics and uncertainty quantification

4.1. Error metrics

In Section 6, we will assess the methods described in Sections 2 and 3 in hindcast experiments against test datasets consisting of sequences of initial data $\{y_0, y_1, \dots, y_{\tilde{N}-1+k}\}$ independent of the training data. If the test data are generated by the same dynamical system $(A, \mathcal{A}, \Phi_t, \alpha)$ as the training data (i.e., in a perfect-model scenario) we have $y_i = \pi(b_i)$, where b_i are states in A such that $b_i = \Phi_{t_i}(b_0)$, $t_i = (i-1)\tau$. In the presence of model error, the y_i are generated by a different dynamical system, $(\tilde{A}, \tilde{\mathcal{A}}, \tilde{\Phi}_t, \tilde{\alpha})$, and we have $y_i = \tilde{\pi}(\tilde{b}_i)$ with $\tilde{b}_i = \tilde{\Phi}_t(\tilde{b}_0)$, $\tilde{b}_i \in \tilde{A}$, and $\tilde{\pi}$ an observation map from \tilde{A} to X .

As stated in Section 2.1, we perform these tests using two distinct types of forecast observables. Observables of the first type are those objectively defined on the test dataset independently of a data analysis algorithm. That is, in the perfect-model case we have the ground-truth values $G_i = f(b_i)$ of the forecast observable f . In the imperfect-model case, the ground-truth values are $G_i = \tilde{f}(\tilde{b}_i)$ where \tilde{f} is the observable map for the dynamical system generating the test data. Examples of objectively defined observables are the components of the state vector, as it will be the case in Section 6.1. Observables of the second type are data-driven observables, constructed through a data analysis algorithm applied to the training data. For example, in Section 6.2, the forecast observables will be the leading low-frequency eigenfunctions of North Pacific SST recovered by cone kernels. For data-driven observables, there is no objective notion of ground truth on the test data, but a natural reference value is given by the out-of-sample extension of f , i.e., $G_i = \hat{F}_0(y_i)$ for the Nyström-based forecast in (2.21), or $G_i = \hat{F}_{0,l}(y_i)$ for the Laplacian pyramids forecast in Section 2.5.2. In all cases, we compute the forecast error at lead time $t = k\tau$ for initial data y_i through the difference $R_t(y_i) = \hat{F}_t(y_i) - G_{i+k}$.

Based on these definitions for the ground truth and the corresponding forecast errors, we compute two time-averaged skill scores, namely the root mean squared error (RMSE) and pattern correlation (PC) scores. These are given by

$$\text{RMSE}(t) = \sqrt{\frac{1}{\tilde{N}} \sum_{i=0}^{\tilde{N}-1} |R_t(y_i)|^2}, \quad \text{PC}(t) = \frac{1}{\tilde{N}} \sum_{i=0}^{\tilde{N}-1} \frac{(\hat{F}_t(y_i) - \bar{F}_t)(G_{i+k} - \bar{G}_k)}{\sigma_F \sigma_G}, \quad (4.1)$$

where $\bar{F}_t = \sum_{i=0}^{\tilde{N}-1} \hat{F}_t(y_i) / \tilde{N}$, $\bar{G}_k = \sum_{i=0}^{\tilde{N}-1} G_{i+k} / \tilde{N}$, $\sigma_F^2 = \sum_{i=0}^{\tilde{N}-1} (\hat{F}_t(y_i) - \bar{F}_t)^2 / \tilde{N}$, and $\sigma_G^2 = \sum_{i=0}^{\tilde{N}-1} (G_{i+k} - \bar{G}_k)^2 / \tilde{N}$.

4.2. Uncertainty quantification

Uncertainty quantification (UQ) is an important part of any statistical forecasting technique, but it appears to have received limited attention in the case of analog forecasting. Here, as a basic form of UQ we estimate the forecast error by local averaging the residual of the training data. In particular, the residual can be evaluated within the training dataset using

$R_t(x_i) = \hat{F}_t(x_i) - f_{i+k}$. Using a normalized local kernel $\hat{\rho}(y, x_i)$ on the training dataset (computed, e.g., via the methods of Section 3), we compute

$$\varepsilon_t^2(y) = \frac{1}{N} \sum_{i=0}^{N-1} \hat{\rho}(y, x_i) |R_t(x_i)|^2, \quad (4.2)$$

and use $\pm \varepsilon_t(y)$ to place two-sided error bars about our forecasts $\hat{F}_t(y)$.

Clearly, this approach is rather rudimentary, and carries a risk of underestimating the forecast uncertainty if \hat{F}_t overfits the training data. However, if steps are taken to prevent overfitting in the training phase (e.g., through the use of the transition probability kernels in (2.30)), it appears that (4.2) provides a reasonable UQ. We defer a more rigorous study of forecast uncertainty in kernel analog forecasting to future work.

5. Comparison with diffusion forecast

Recently, Berry et al. [7] developed a technique for nonparametric forecasting of dynamical systems called diffusion forecasting, which has some common aspects with the kernel analog forecasting methods developed here. In this Section, we discuss the similarities and differences of these techniques and comment on their relative strengths and weaknesses.

Similarly to kernel analog forecasts, diffusion forecasts do not assume a parametric structure for the forecast model, and rely instead on observations of the system state in the past (analogous to the training samples $\{x_i\}$ used here) to approximate operators (analogous to U_t) governing the evolution of observables. Moreover, as with kernel analog forecasting, kernel techniques also play an important role in diffusion forecasts. There, the eigenfunctions from a variable-bandwidth kernel [8] are used to build a data-driven orthonormal basis $\{\phi_0, \phi_1, \dots\}$ for the L^2 space of the dynamical system, and this basis is used to create a matrix representation of the forward operator of the dynamics. More specifically, the dynamical system in diffusion forecasting is assumed to evolve on a manifold M , and the $\{\phi_i\}$ basis is orthonormal on $L^2(M, \mu)$ space associated with the system's invariant measure μ . Diffusion forecasting can handle both deterministic and stochastic systems; in the former case it approximates the Koopman operator, and in the latter case the Kolmogorov operator. These operators are represented by $l \times l$ matrices with elements $A(t)_{ij} = \langle \phi_i, U_t \phi_j \rangle$. Given an observable $f = \sum_{i=0}^{l-1} c_i \phi_i$, and an initial probability measure with density $\rho = \sum_{i=0}^{l-1} b_i \phi_i \in L^2(M, \mu)$ relative to μ , the diffusion forecast at lead time t is given by the expectation value

$$\int_M U_t f \rho d\mu \approx \sum_{i,j=0}^{l-1} A(t)_{ij} b_i c_j.$$

An advantage of diffusion forecasting over kernel analog forecasting is its greater flexibility with regards to the initial probability measure. In kernel analog forecasting, the initial probability density is a kernel-dependent function of the initial data which is

cumbersome to explicitly control. On the other hand, the initial density in diffusion forecasts is entirely up to the user, provided of course that it can be well represented in the eigenfunction basis. Another advantage of diffusion forecasting is that it is applicable to both deterministic and stochastic systems. It is plausible that kernel analog forecasting can be generalized to the stochastic case using similar error estimation techniques as in [7].

An advantage of kernel analog forecasting over diffusion forecasting is its ability to handle partially observed systems. Applying diffusion forecasting to the partially observed systems studied here would produce an operator of the form $\Pi_l U_t \Pi_l$ where Π_l is the orthogonal projector to the l -dimensional subspace of $L^2(A, \alpha)$ spanned by $\{\phi_0 \circ \pi, \phi_1 \circ \pi, \dots, \phi_{l-1} \circ \pi\}$. Unless that subspace happens to be an invariant subspace of the dynamics, forecasting with $\Pi_l U_t \Pi_l$ will lead to biases. The analog forecasting with Nyström extension in Section 2.4 suffers from similar biases, but the techniques in Sections 2.2 and 2.5 which are based on averaging operators will produce more uncertain, rather than biased, forecasts when faced with partial observations. Kernel analog forecasting also appears to make fewer assumptions than diffusion forecasting about the smoothness of the dynamical system and the manifold, though it is plausible that some of the smoothness assumptions in diffusion forecasting could be relaxed. In particular, the theoretical development of diffusion forecasting in [7] is for dynamical systems on smooth manifolds, but it is possible that a similar framework could be developed for more general measurable spaces (in fact, in [7] it was shown that diffusion forecasting can be successfully applied to the Lorenz 63 system, where the attractor clearly violates the smooth manifold assumption).

Overall, while kernel analog forecasting and diffusion forecasting have their individual strengths and weaknesses, both methods are useful in a variety of forecasting scenarios where skillful parametric models are hard to construct. Even in situations where parametric models are available, these techniques can provide useful benchmarks for the forecast skill that should at a minimum be expected from parametric models.

6. Experimental results

6.1. A chaotic intermittent low-order atmosphere model

We illustrate several key aspects of the forecasting methods described in Sections 2 and 3, including kernel analog forecasting with Laplacian pyramidseq:LapApprox, lagged sequences in (3.1), and cone kernels in (3.5), through experiments on a six-dimensional low-order model of the atmosphere featuring chaotic metastability [?, 23]. Specifically, setting $u(t) = (u_1(t), \dots, u_6(t)) \in \mathbb{R}^6$, we consider a deterministic dynamical system of the form $\dot{u}(t) = V(u(t))$, where the vector field, $V(u) = (V_1(u), \dots, V_6(u))$, is given by

$$\begin{aligned} V_1(u) &= \gamma_1^* u_3 - C(u_1 - u_1^*), & V_2(u) &= -(\alpha_1 u_1 - \beta_1) u_3 - C u_2 - \delta_1 u_4 u_6, \\ V_3(u) &= (\alpha_1 u_1 - \beta_1) u_2 - \gamma_1 u_1 - C u_3 + \delta_1 u_4 u_5, & V_4(u) &= \gamma_2^* u_6 - C(u_4 - u_4^*) + \epsilon(u_2 u_6 - u_3 u_5), \\ V_5(u) &= -(\alpha_2 u_1 - \beta_2) u_6 - C u_5 - \delta_2 u_4 u_3, & V_6(u) &= (\alpha_2 u_1 - \beta_2) u_5 - \gamma_2 u_4 - C u_6 + \delta_2 u_4 u_2. \end{aligned}$$




Figure 6.1: Scatter plot of the state vector components (u_1, u_4) from a 10,000-day simulation of the low-order atmospheric model, showing approximate locations of the zonal and blocked regimes.

Here, α_i , β_i , γ_i , γ_i^* , ϵ , C , and u^* are model parameters. This model is derived by a low-order truncation of the streamfunction in the barotropic vorticity equation in a channel in the presence of topography and a zonal (east–west) forcing profile. In particular, the state-vector components u_i are expansion coefficients in the spatial Fourier basis for the channel. Physically, this model provides a coarse approximation of atmospheric dynamics in the presence of orography and a zonal background flow—it is a simplified model for the polar jet stream interacting with continental mountain ranges, such as the mountain ranges of North America. Observationally, it is known that such flows exhibit metastable transitions between so-called zonal and blocked states, where the jet propagates relatively unimpeded, or is significantly blocked by the presence of topography, respectively. With an appropriate choice of parameters [23], the six-dimensional model exhibits a qualitatively similar behavior, and the two regimes are manifested by distinct regions in the (u_1, u_4) plane (see Figure 6.1). In the spatial domain, these components of the state vector represent purely zonal flow, with no variation in the meridional direction, reflecting the interesting regime switching between zonal and blocked states.

In what follows, we apply the techniques of Sections 2 and 3 to predict the evolution of the observable $f = (u_1, u_4)$ representing the regimes, using as initial data either the full state vector, $y = u$, or a vector formed by applying delay-coordinate maps to the projected state vector $z = (u_1, u_4)$ in accordance with (3.1). We note that forecasting with reduced parametric models is particularly challenging in this example. For instance, in [23] it was found that even a modest dimension reduction of the full model to four degrees of freedom by Galerkin projection onto the PCA basis produces models that either decay to fixed points, or become locked on orbits of unrealistically high temporal regularity, failing to reproduce chaotic regime transitions. On the other hand, Galerkin projection onto a basis derived

by eigenfunctions from the kernel in (3.5) with $\zeta = 0$ was able to exhibit chaotic regime transitions using only three dimensions [32]. This suggests that the patterns derived from Laplace-Beltrami eigenfunctions are more intrinsic to the dynamical system than the linear PCA modes.

Our training dataset, \hat{M} , consists of $N = 10^4$ samples taken every $\tau = 1$ day (d) from an equilibrated simulation of the full dynamical system performed using a Runge-Kutta method. A scatter plot of the components (u_1, u_4) of the state vector in the training data is depicted in Figure 6.1. To generate the test dataset, we changed the initial conditions and ran another equilibrated simulation with $\tilde{N} = 5000$ data points.

First, we discuss hindcast experiments using the full state vector as initial data. In these experiments we did not perform delay-coordinate mapping as the initial data are complete [i.e., we set $q = 1$ in (3.1)]. Therefore, the only dependence of the kernel (3.5) in the dynamics is through the finite-difference estimates, v and w , of the dynamical vector field V , which were evaluated here using a first-order backward scheme. Figure 6.2 shows sample forecast trajectories and the RMSE and PC skill scores from (4.1) obtained via the empirical forecasting methods described in Sections 2 and 3; namely, conventional analog forecasting with Euclidean distances, single analog forecasting with dynamics-adapted kernel affinity, kernel analog forecasting with Laplacian pyramids, and constructed analog method introduced in [62]. The constructed analog method makes prediction based on a linear combination of the historical data, and the weights for averaging are determined by fitting the initial data with least squares. Note that we do not use the Nyström approach in (2.21) as f is not a bandlimited observable in the eigenfunction basis.

For the affinity-based analog forecasts, we used the cone kernel in (3.5) with $\zeta = 0.995$, and also examined the case with $\zeta = 0$ corresponding to no influence of the angular terms. We find that the analogs chosen with respect to the $\zeta = 0$ kernel affinities via (2.26) have essentially identical skill to the conventional analog forecast, but there is a marked improvement for the cone kernel with $\zeta = 0.995$. The improvement of skill is reflected in the improved RMSE and PC results in Figures 6.2(b) and 6.2(c), and is more prominent in the forecast time series in Figure 6.2(a). There, the analog based on Euclidean distance deviates from the ground truth for $u_1(t)$ after $t \approx 360$ d, but the affinity-based analog with $\zeta = 0.995$ tracks the true trajectory up to $t \approx 480$ d. We therefore see that choosing analogs with respect to an affinity measure that preferentially selects samples with similar time tendency (as $\zeta \approx 1$ cone kernels do) leads to improvement of skill. For constructed analog method, finding the linear combination weights is challenging. Since we have more samples in time than the number of spatial points, we need to solve an under-determined system to get the weights. The least squares solution might not be the optimal solution. Therefore, the forecast skill of the constructed analogue method is low compared to other analog forecasting methods and the sample trajectory is not shown in Figure 6.2a.

The most significant improvement of skill takes place for the kernel analog forecasting with Laplacian pyramids in (2.26). Here, we used $\zeta = 0.995$, and truncated the kernel \hat{p}_{ϵ_l}

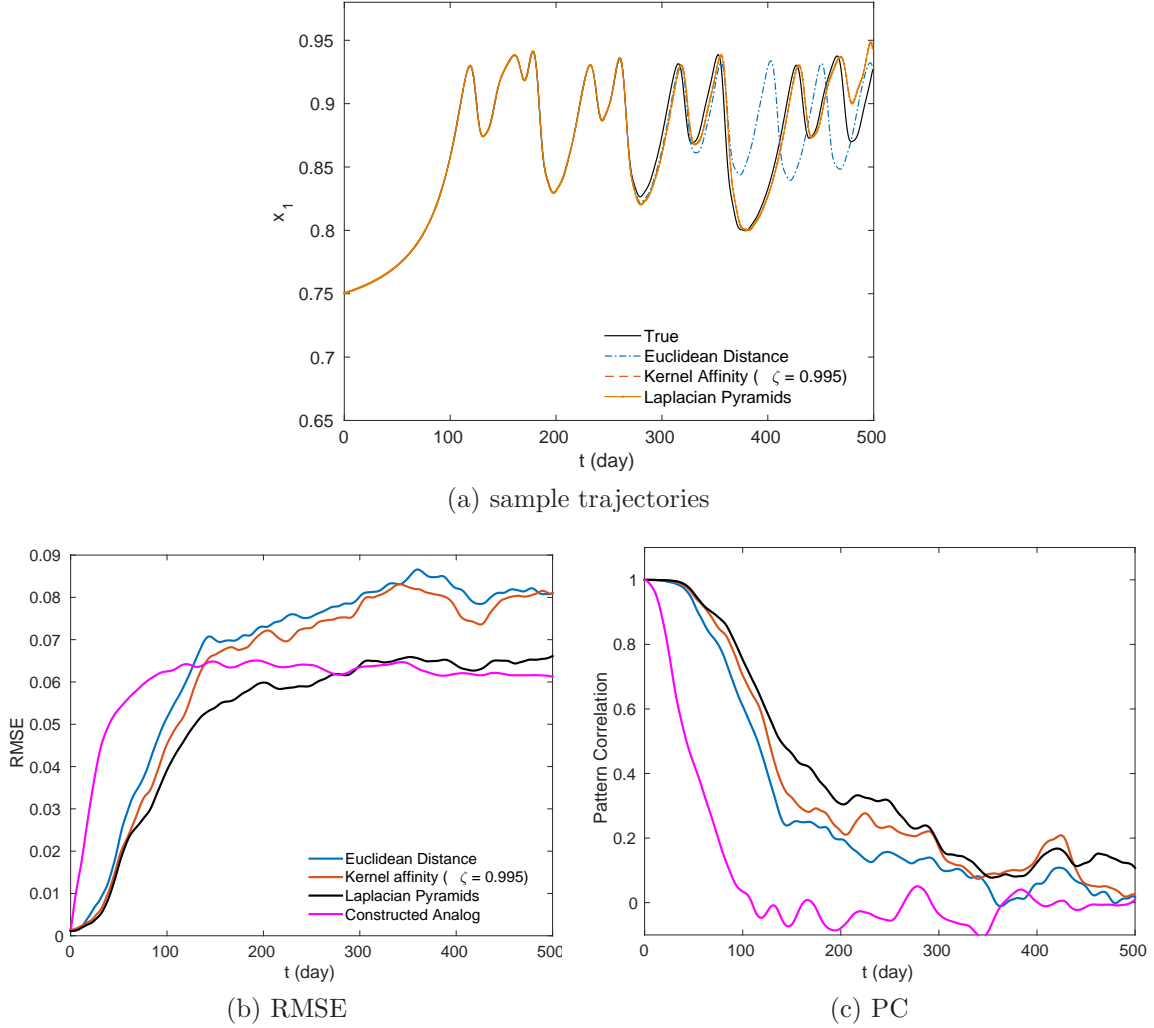


Figure 6.2: Hindcast experiments for the low-order atmospheric model with complete initial data. (a) Ground-truth and predicted time series for u_1 via conventional Euclidean-distance analogs, Kernel affinity-based analogs with $\zeta = 0.995$, and kernel analog forecasting with Laplacian pyramids. (b,c) Root mean squared error (RMSE) and pattern correlation (PC) scores, respectively, for the forecast models in (a).

with 10 nearest neighbors. The initial bandwidth parameter ϵ_0 was set to the median of the distances between the 10 nearest neighbors. The iteration of the Laplacian pyramids stopped at $l = 2$ iterations. In Figure 6.2, the kernel analog forecasting accurately tracks the true trajectory to at least 480 d, the RMSE is lower at both intermediate and long times ($t \gtrsim 60$ d), and the PC values exceed those from other methods for up to 300 d leads. To summarize, in these experiments the cone-kernel geometry with $\zeta \approx 1$ (which produces neighborhoods aligned with the dynamical flow), improves the identification of single analogs, and moreover there are further benefits from using multiple samples in the training data through Laplacian pyramids.

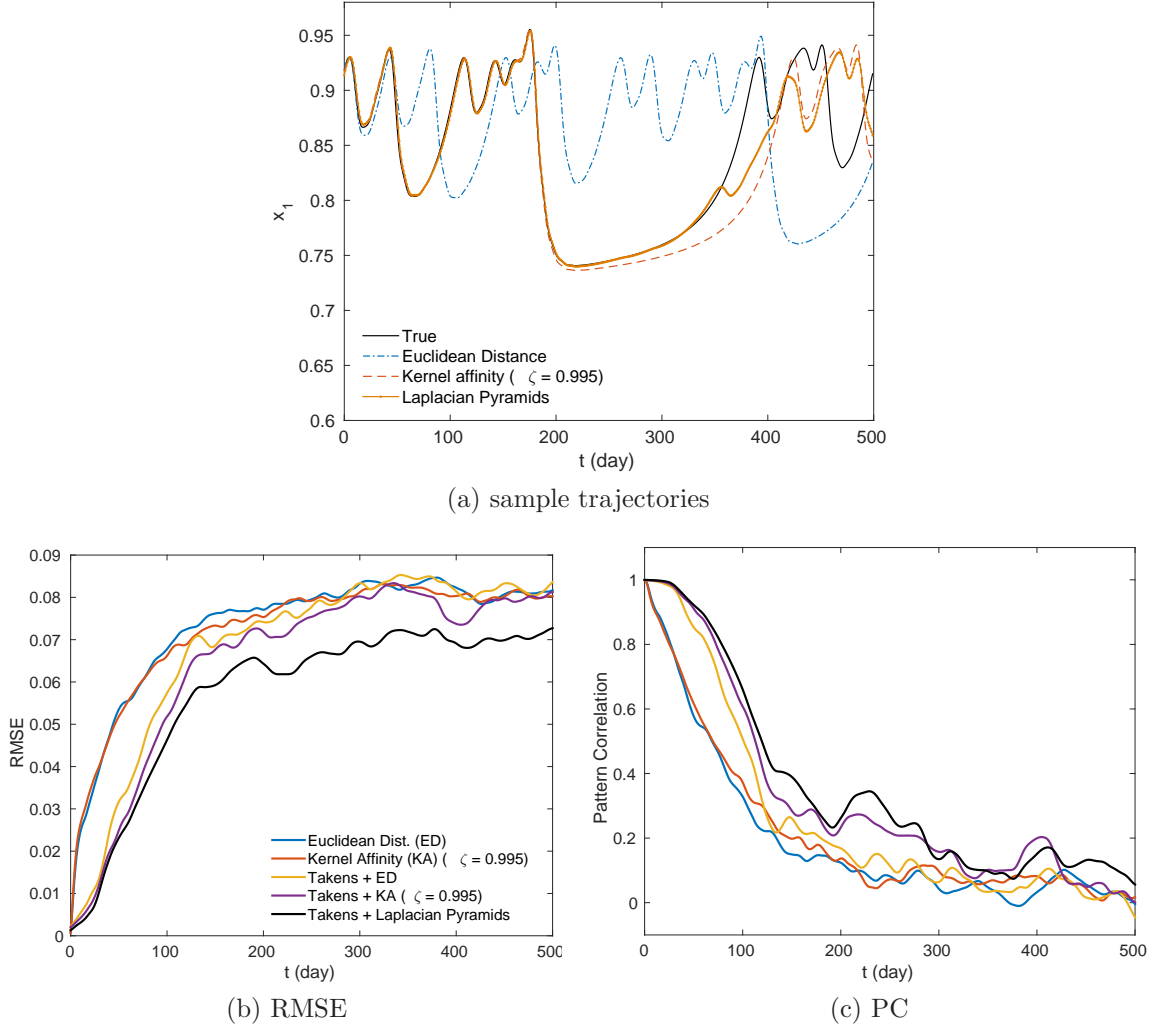


Figure 6.3: Same experiments as in Figure 6.2, but for incomplete initial data, $z = (u_1, u_4)$, reconstructed using the delay-coordinate mapping in (3.1) with q lags. The experiments shown here are for $q = 1$ and $q = 22$, corresponding to no delay-coordinate mapping and delay coordinate mapping over a physical time window of 21 days, respectively.

Next, to mimic a real-world scenario where only partial observations are available, we perform an experiment where only two coordinates, $z = (u_1, u_4)$, of the full state vector are observed. Following 3.1, we form the vectors of initial data x_i by applying delay-coordinate mapping to z_i with q lags. In an operational setting, this would correspond to predicting the state of the atmosphere given its observed evolution over q days in the past. The results from hindcast experiments with $q = 1$ and $q = 22$ [i.e., no delay coordinate mapping and delay coordinate mapping over a physical time window $(q - 1)t = 21$ d, respectively] and the same forecasting techniques used in the fully-observed case are shown in Figure 6.3.

Without Takens' time-lagged embedding, the analog forecasts with partial observations perform significantly worse than the same analog forecasts with complete observations for

all forecasting methods. This is consistent with the two-dimensional initial data with $q = 1$ being insufficient to determine neighborhoods on the attractor, and there is evidence that the intrinsic dimension of the attractor in this class of systems is greater than two [23]. On the other hand, using Takens’ time-lagged embedding with $q = 22$ delays leads to nearly equal skill as the case with full observations. Thus, in this relatively low-dimensional example, delay-coordinate mapping is a highly effective way of recovering the necessary dynamical information for empirical forecasting. Similarly to the complete-observation experiments, the $\zeta = 0.995$ cone-kernel affinity measure gave the best single-analog forecast, and the kernel analog forecasting with Laplacian pyramids provides the highest skill across the board.

6.2. Long-range forecasting in the North Pacific sector of a comprehensive climate model

As a high-dimensional application, we study low-frequency SST variability in the North Pacific sector of the coupled climate model CCSM3 [19]. The datasets used in this study are two long equilibrated control integrations available in the public domain (<https://www.earthsystemgrid.org>) with designations b30.004 and b30.009. Integration b30.004 spans 900 years (y) using a 2.8° nominal resolution for the atmospheric component of the model; b30.009 spans 500 years, but employs a higher, 1.4° , atmosphere resolution. Both integrations employ the same 1° ocean resolutions. Two sets of experiments are discussed below, the first of which is a perfect-model experiment using the first 400 years of integration b30.004 as the training dataset and the second half of b30.004 as the test dataset. The second set of experiments introduces model error and we use the data from b30.009 as test data from “nature” and the first 800 years of b30.004 as training. The training datasets are also used to construct data-driven low-frequency observables for prediction. Elsewhere [20], we perform hindcast experiments against actual observational data acquired via remote sensing, involving multivariate datasets of SST, regional arctic sea ice concentration and sea ice volume. The b30.004 and b30.009 datasets were also used in [31], where the dominant spatiotemporal modes of North Pacific SST variability as extracted by NLSA algorithms were compared across different GCMs.

6.2.1. Recovering low-frequency observables. Following [31, 30], we extract low-frequency SST modes using $q = 24$ delay-coordinate lags (corresponding to a two-year temporal embedding window) to induce timescale separation. Throughout, we work with cone kernels with $\zeta = 0.995$ as these kernels provide better timescale separation in the eigenfunctions [30] than the $\zeta = 0$ kernels used in [31]. The SST snapshots have dimension $d = 6671$ (equal to the number of North Pacific ocean gridpoints in CCSM3), and the dimension of the data vectors x_i after delay-coordinate mapping via (3.1) is $m = 24 \times 6671 = 160,104$. Throughout, we work with a kernel bandwidth parameter $\epsilon = 2.25$, but our results are qualitatively robust for bandwidths in the interval 0.5–5.

Representative eigenfunctions computed from the training data with these delay-coordinate and kernel parameters are shown in Figure 6.4. The eigenfunctions separate




Figure 6.4: Representative eigenfunctions of North Pacific SST in the CCSM3 model. The eigenfunctions, $\phi_i(x_j)$, are shown as time series with respect to the timestamps $t_j = (j-1)\tau$ for a 40-year segment of the 400-year training dataset. Frequency spectra, computed by applying the discrete Fourier transform of the 400-year time series are also shown. (a) Annual periodic mode; (b) $L_1 = \phi_5$, first low frequency mode, Pacific decadal oscillation (PDO); (c) annual intermittent mode corresponding to the PDO; (d) $L_2 = \phi_8$, second low frequency mode.

the temporal variability of the data into qualitatively distinct families, namely periodic [Figure 6.4(a)], low-frequency [Figures 6.4(b,d)], and intermittent [Figure 6.4(c)] modes. The periodic modes closely resemble sinusoids with frequencies given by integer multiples of 1 y^{-1} . These modes represent the harmonics of the seasonal cycle (a prominent source of variability in North Pacific SST), and form doubly-degenerate pairs with a phase offset of $\pi/2$. The low-frequency modes describe oscillations taking place at interannual to decadal timescales, and are characterized by red-noise-like power spectra. The intermittent modes consist of periodic signals at the seasonal-cycle harmonics modulated by low-frequency envelopes. We refer the interested reader to [31, 14, 30] for detailed discussions on the properties and physical significance of these mode families.

In what follows, our focus will be on forecasts of the leading-two low-frequency eigenfunctions, i.e., $L_1 = \phi_5$ and $L_2 = \phi_8$. These eigenfunctions carry significant power

on interannual to decadal timescales, making them good candidate observables for long-range forecasting, but in addition they are associated with physically meaningful patterns in spatiotemporal reconstructions (e.g., see the online supporting animation in [30]). In particular, L_1 and L_2 display the salient features of two prominent low-frequency climate patterns in North Pacific. The first low frequency mode represent PDO [43] and the second mode is a manifestation of ENSO, which is most prominent in tropical Pacific [61], respectively. The former is characterized by a horseshoe-like temperature anomaly pattern, developing east of Japan, together with an anomaly of the opposite sign along the west coast of North America. Based on these observations, we select eigenfunctions L_1 and L_2 as data-driven observables for empirical forecasting. It is important to note that our kernel construction is essential to achieve timescale separation; in particular, if no Takens embedding is used ($q = 1$) the temporal character of the eigenfunctions becomes corrupted, mixing low-frequency and periodic variability. However, despite their low-frequency character, L_1 and L_2 are particularly challenging to predict with parametric models, and first-order autoregressive models often fail to beat the persistence forecast [20]. In the predictive skill results discussed below we include the persistence forecast as a reference for the forecast skill that can be achieved in these time series with simple autoregressive models. We used constructed analog method to predict the evolution of the leading two low-frequency modes. The monthly climatology is removed from both the training and test datasets to find the coefficients for the weighted averaging. However, the pattern correlation scores for this method on the first two low-frequency modes are below 0.5 at $t = 0$. Therefore, these results are not included in the following subsections.

6.2.2. Forecasting in a perfect-model environment. We begin with an application of the techniques of Sections 2 and 3 in hindcasts of a test dataset generated by the same model as the model generating the training data—as stated earlier, the test dataset in this set of experiments consists of the 400 years of CCSM3 integration b30.004 following the 400 years of the training data. Figure 6.6 display predictive skill results for L_1 and L_2 obtained with the methods described in Section 2, as well as with the persistence forecast. The prediction observables in these experiments are pure eigenfunctions, so we performed kernel analog forecasting using both the Nyström method and Laplacian pyramids via (2.21) and (2.26), respectively. Moreover, because the prediction observables are data driven (i.e., they are not objectively defined on the test dataset), the skill scores were evaluated treating the Nyström out-of-sample extension of L_1 and L_2 (Figure 6.6) as ground truth. The persistent forecast follows the Nyström out-of sample extension of the functions at lead time $t = 0$.

Using Laplacian pyramids for predicting the first two low frequency modes L_1 and L_2 , we chose the initial kernel bandwidth to be the median of all pairwise distances. The training error in Laplacian pyramids first decreases with the increasing number of iterations, and then increases (see Figure 6.5). The iterations stop at the minimum training error. For this dataset, 8 or 9 iterations are needed. Because of the choice of the ground truth for the




Figure 6.5: Training errors with respect to the number of iterations in Laplacian pyramids for $L_1 = \phi_5$ at different lead time t . The initial bandwidth ϵ_0 is chosen to be the median of all pairwise distances.

computation of error metrics, the Laplacian pyramids estimation has small initial error at $t = 0$.

The main finding from these experiments is that the kernel analog forecasting, implemented either with the Nyström extension in (2.21) or with Laplacian pyramids in (2.26), are the best performers across the board, substantially improving over single-analog forecasting approaches as well as persistence forecasts. In the case of the PDO, PC scores exceeding 0.5 (a popular lower threshold for “useful” forecasts) persist for 32-month lead times using Nyström method and 30-month lead times for Laplacian pyramids. In contrast, the persistence forecast crosses the $PC = 0.5$ threshold at 25-month leads, and has decayed to 0.38 at $t = 32$ months. Note that for this dataset the persistence forecast performs better than linear autoregressive models [20], which are popular parametric models for the PDO [46]. Improvements over the persistence forecast are even more significant for the second low frequency mode L_2 , which is a more rapidly decorrelating variable than the PDO. In this case, the $PC = 0.5$ crossing time for the kernel-weighted forecasts is 22 months, which is more than a factor of two improvement compared to the 9-month crossing time for the persistence forecast.

Consider now the forecasts made by single-analog methods. In this example with a high-dimensional ambient data space, the analogs selected via the cone-kernel affinity measure using (3.5) have nearly identical skill as the conventional Euclidean-distance analogs (cf. the low-dimensional example in Section 6.1). Therefore, the affinity-based results are not included in Figures 6.6. Note that in the literature it is customary to calibrate the forecast error vs. lead time curves such that the error at $\tau = 0$ vanishes (e.g., [12]), and our skill scores with single analogs will generally appear lower in those approaches.

The prediction results of kernel analog forecast with Nyström extension for L_1 and L_2

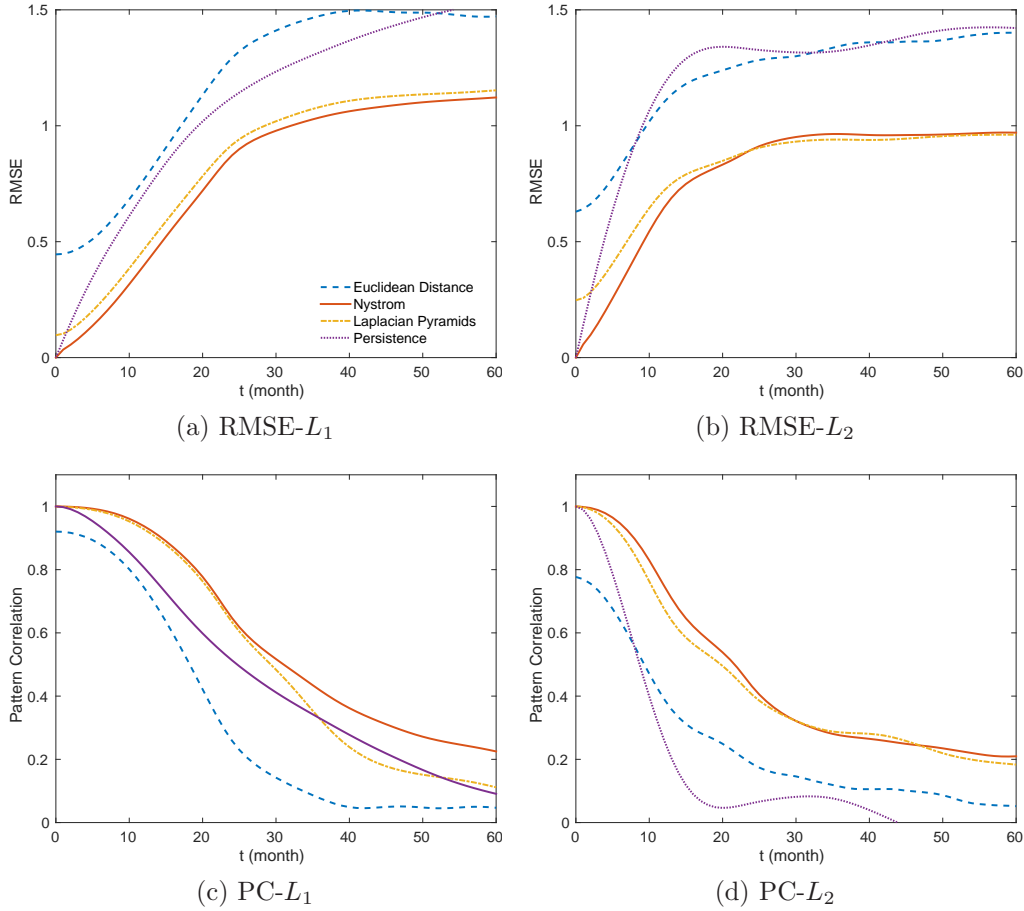


Figure 6.6: Forecast skill for the low-frequency eigenfunctions $L_1 = \phi_5$ and $L_2 = \phi_8$ in a perfect-model scenario. The RMSE and PC scores from (4.1) were evaluated for the persistence forecast, single-analog forecast based on Euclidean distances, and kernel analog forecasting with Nyström extension and Laplacian pyramids.

at 5, 10, 20 months lead are illustrated in Figure 6.7. At 5 months lead, the predicted time series agree very well with the true signals. The estimated uncertainty is also very small. When lead time t increases, the prediction error increases and the uncertainty ε_t in (4.2) becomes larger (see the gray shaded area in Figure 6.7).

It is evident from Figure 6.6 that the single-analog method is affected by large error at short lead times. This error is caused by poor reconstruction accuracy using a single analog. That is, unless the sampling density on the attractor is large, the ground truth at initialization time ($t = 0$) may deviate significantly from the estimated value $\hat{F}(y, 0)$ from the analog. Substantial errors at initialization time typically persist at least for the short term, and in such cases single-analog methods will have lower short-term skill than methods which are able to accurately reconstruct the forecast observable at initialization, including the persistence forecast. Indeed, as shown in Figures 6.6(a,c), the single-analog forecast for

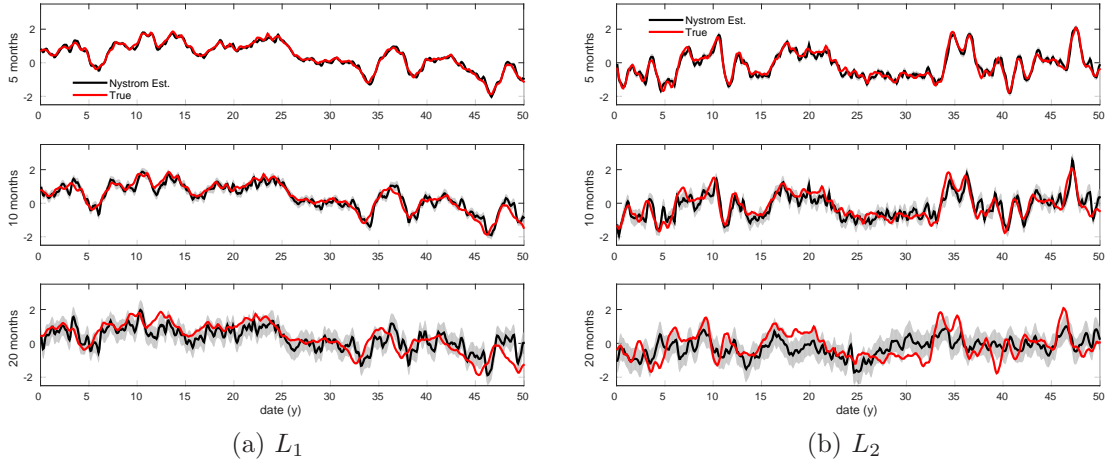


Figure 6.7: Prediction of (a) L_1 and (b) L_2 at 5, 10, and 20 months lead. The red line shows the true signal and the black line shows the kernel analog forecasting with Nyström extension. The gray shaded areas indicate uncertainty with $\pm\epsilon_t$ defined in (4.2).

L_1 (a slowly-decorrelating observable) fail to beat persistence for both short and long times. Single-analog forecast of L_2 , which is a more rapidly decorrelating variable, does outperform persistence at moderate to long lead times [$t \gtrsim 10$ months in Figures 6.6(b,d)], but the short-term skill for this observable is also poor. Note that poor short-term skill with single-analog method was not an issue in the experiments in Section 6.1 with the low-order atmospheric model (see Figures 6.2 and 6.3). There, the dimension of both the ambient data space and the attractor was low, and the sampling density was sufficiently for accurate reconstruction at short leads even with a single analog. In contrast, by appropriately weighing multiple samples in the training data, the kernel analog forecasting performs well in both low- and high- dimensional ambient spaces, and for both short and long lead times.

The computational complexity for Nyström extension is $O(N^3 + l_1 N)$, where the first term, $O(N^3)$, is the complexity for computing the eigendecomposition of the $N \times N$ kernel matrix and the second term is for extending the eigenfunctions and reconstructing the signal with the top l_1 eigenfunctions. The computational complexity for Laplacian pyramids is $O(l_2 N^2)$, where l_2 is the number of iterations. In this example, the eigendecomposition of the appropriate kernel matrix was performed to derive the data-driven observables, L_1 and L_2 , and therefore, the kernel analog forecasting with Nyström extension only takes $O(l_1 N)$ operations. The running time is 85 seconds for Nyström extension and 276 seconds for Laplacian pyramids implemented with MATLAB on a machine 16 cores with 500 GB RAM. Although Laplacian pyramids is slower than Nyström extension, in a more general problem, where the observables are not Laplace-Beltrami eigenfunctions, it can be more efficient than Nyström method.

6.2.3. Forecasting with model error in the training data To test the robustness of our techniques in the presence of model error in the training data, we now study the forecast skill of the first two low frequency modes L_1 and L_2 in an experiment where the training data come from the first 800 years of the b30.004 CCSM integration, but the test data come from 500 years of the b30.009 integration. As mentioned earlier, these two integrations differ in the resolution of the atmospheric component of CCSM. These two datasets therefore have different atmospheric dynamics, which lead in turn to differences in SST variability through the associated atmosphere–ocean momentum and heat fluxes. Because b30.009 has higher atmospheric resolution than b30.004, in these experiments we view b30.009 as “nature” and b30.004 as an imperfect model. More broadly, this experimental setup is representative of many empirical modeling applications in the physical sciences, where long-term observations of nature are not available for training.

The skill scores in Figure 6.8 were evaluated treating the Nyström out-of-sample extension of L_1 and L_2 from b30.004 to b30.009 as ground truth, since the prediction observables are pure eigenfunctions. The persistence forecast follows the Nyström out-of-sample extension at time $t = 0$. As expected, the skill scores in these experiments with model error in the training data are generally lower than their perfect-model counterpart in Figures 6.6. In the case of the first low frequency eigenfunction L_1 , PC scores exceeding 0.5 persist for 30-month lead times, i.e., 2 months less than in the perfect model experiment. In contrast, the persistence forecast crosses the $PC = 0.5$ threshold at 24-month leads, and has decayed to 0.40 at $\tau = 30$ months. The kernel analog forecasting methods still provide significantly higher skill than persistence, and also outperform single-analog forecasting using Euclidean distances. In the case of the second low frequency mode L_2 , the $PC = 0.5$ crossing time for the kernel analog forecasts remains more than a factor of two higher than the persistence forecast [see Figure 6.8(d)]. Figure 6.9 shows the Nyström kernel analog forecasting results for L_1 and L_2 at 5, 10, and 20 months lead time.

7. Conclusions

In this paper, we proposed a family of data-driven, nonparametric forecasting methods for observables of dynamical systems. Inspired by Lorenz’s analog forecasting technique [41], these methods perform forecasting through local averaging operators that predict the value of the forecast observable by a weighted average over multiple states in the training data that closely resemble the initial data. These operators are constructed through suitably normalized kernels which measure the pairwise similarity of data taking dynamics into account. In particular, the kernels used here incorporate empirically accessible aspects of the dynamical system through Takens delay-coordinate maps [32, 6] and finite-difference approximations of the dynamical vector field on the attractor [30], enhancing timescale separation and robustness to changes of observation modality. Due to these features, the properties of the associated averaging operators can be understood in terms of a modified Riemannian geometry favoring samples in the training data with similar dynamical evolution

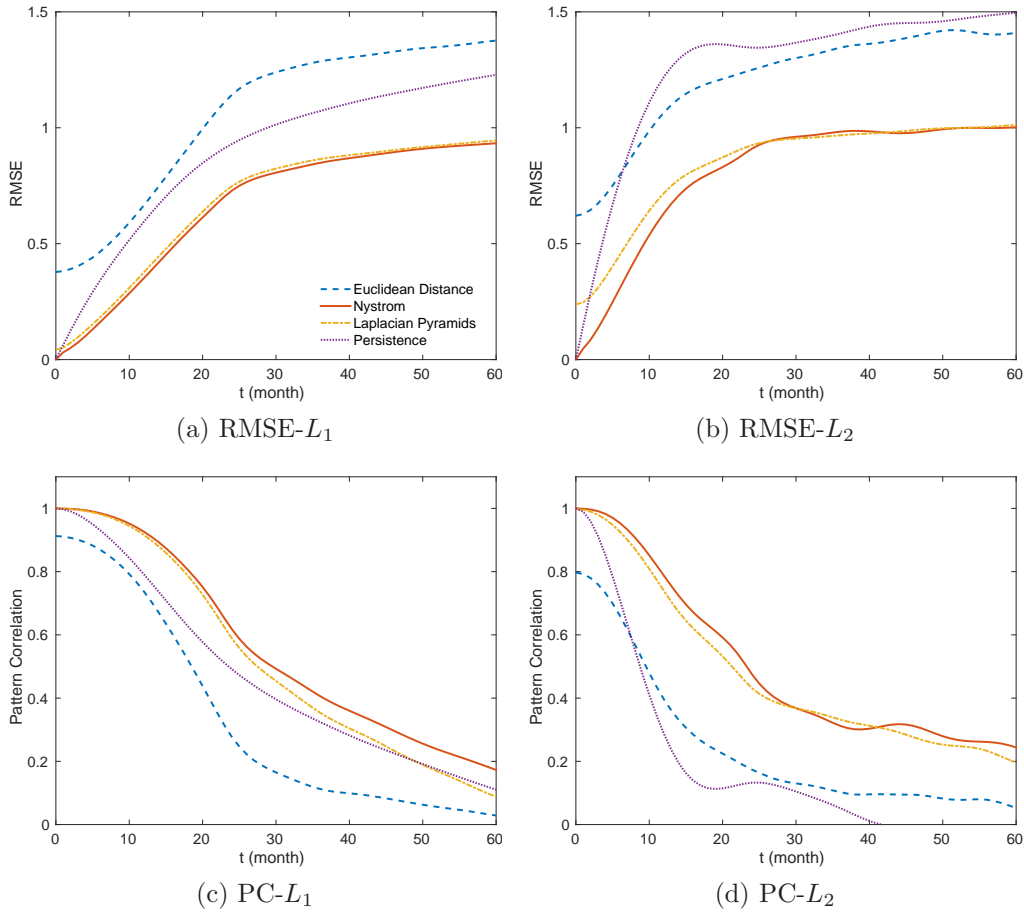


Figure 6.8: Forecast skill for the low-frequency eigenfunctions $L_1 = \phi_5$ and $L_2 = \phi_8$ respectively, for training data with model error. Similarly to Figure 6.6, the RMSE and PC scores were computed for the persistence forecast, single-analog forecast based on Euclidean distances, and kernel analog forecasting using the Nyström extension and Laplacian pyramids.

to the observed data at forecast initialization.

Mathematically, kernel analog forecasting has strong connections with kernel methods for out-of-sample extension of functions. Here, we studied two such approaches based on Nyström extension [18] and the Laplacian pyramids technique [50, 28]. The Nyström-based approach naturally lends itself to the identification and forecasting of data-driven observables constructed from kernel eigenfunctions with high smoothness and favorable predictability. On the other hand, kernel analog forecasting with Laplacian pyramids performs well with more general classes of forecast observables. Used in conjunction with a bistochastic kernel normalization [16], this method is also able to make unbiased forecasts from incomplete initial data.

We demonstrated the efficacy of these schemes in hindcast experiments involving a

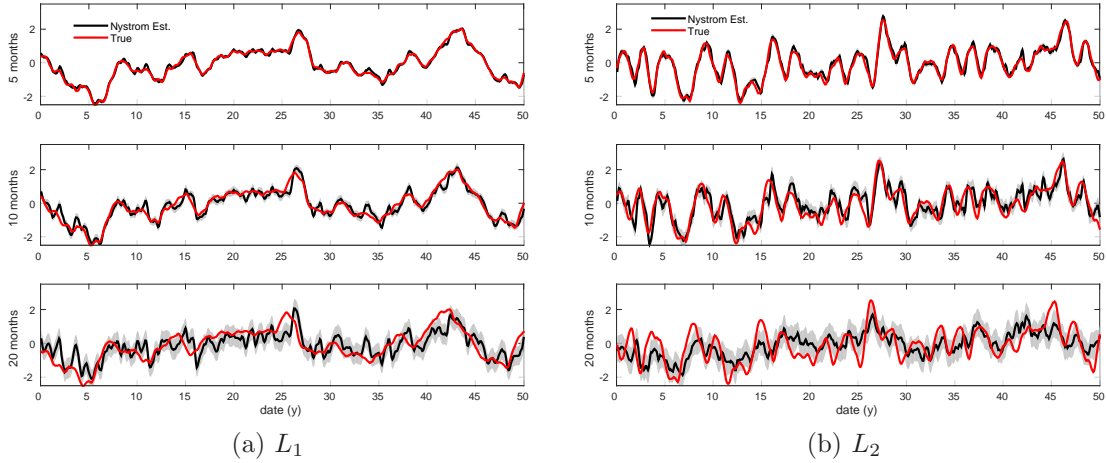


Figure 6.9: Prediction of (a) L_1 and (b) L_2 with model error at 5, 10, and 20 months lead. The red line shows the true signal and the black line shows the kernel analog forecasting with Nyström extension. Light gray shaded area indicate uncertainty with $\pm\epsilon$ in defined in (4.2).

low-order deterministic model for the atmosphere with chaotic regime metastability, as well as interannual to decadal variability of sea surface temperature (SST) in the North Pacific sector of a comprehensive climate model. We also studied a North Pacific SST experiment with model error in the training data. In all experiments, our proposed weighted-ensemble methods led to significant improvement of forecast skill compared to conventional analog forecasting with Euclidean distances. The kernel analog forecasting also significantly outperformed forecasts made with single analogs identified using Euclidean distances.

In the North Pacific SST applications, we constructed data-driven forecast observables L_1 and L_2 through kernel eigenfunctions. The timescale separation capabilities of the kernels used here contributed significantly to the clean low-frequency character and favorable predictability properties of the observables, extending the range of useful forecasts [> 0.5 pattern correlation score (PC)] to 2.6 years for the PDO and 1.8 years for the the second low frequency eigenfunction that is associated with ENSO. As a benchmark, the persistence forecast for these variables (which is more skillful than simple autoregressive models) decays below the 0.5 PC threshold at 2 years (L_1) and 0.8 years (L_2). Again, kernel analog forecasting with Nyström extension and Laplacian pyramids (as opposed to single analog) were crucial to obtain these results.

Stemming from this work are several future research directions. Here, we focused on forecasting with the kernels developed in [32, 34], but other kernels for dynamical systems have been proposed in the literature [54, 58, 8] and could be employed in our forecasting schemes. Moreover, we have focused on Nyström-type extensions and Laplacian pyramids, but other interpolation and extrapolation techniques, such as kriging, can be explored.

Finally, we have studied the asymptotic behavior of our algorithms in the limit of large data, but have not addressed the rate of convergence to the asymptotic limit. In particular, it should be fruitful to study the convergence properties of our methods in different classes of dynamical systems (e.g., geometrically ergodic systems), and related implications to risk of overfitting and uncertainty quantification.

Acknowledgments

D. Giannakis and Z. Zhao wish to acknowledge support from ONR MURI Grant 25-74200-F7112 and NSF Grant DMS—1521775. The research of D. Giannakis is partially supported by ONR DRI grant N00014-14-0150. The authors wish to thank Darin Comeau and Andrew Majda for stimulating conversations.

References

- [1] N. Aronszajn. Theory of reproducing kernels. *Trans. Amer. Math. Soc.*, 68(3):337–404, 1950.
- [2] N. Aubry, W.-Y. Lian, and E. S. Titi. Preserving symmetries in the proper orthogonal decomposition. *SIAM J. Sci. Comput.*, 14:483–505, 1993.
- [3] M. Belkin and P. Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.*, 15:1373, 2003.
- [4] M. Belkin and P. Niyogi. Convergence of Laplacian eigenmaps. In *Advances in Neural Information Processing Systems*, pages 129–136. MIT Press, 2007.
- [5] Y. Bengio, J.-F. Paiement, P. Vincent, O. Delalleau, N. L. Roux, and M. Ouimet. Out-of-sample extensions for LLE, isomap, MDS, eigenmaps, and spectral clustering. In *Advances in Neural Information Processing Systems*, pages 177–184. MIT Press, 2003.
- [6] B. Berry, R. Cressman, Z. Gregurić Ferenč, and T. Sauer. Time-scale separation from diffusion-mapped delay coordinates. *SIAM J. Appl. Dyn. sys.*, 12:618–649, 2013.
- [7] T. Berry, D. Giannakis, and J. Harlim. Nonparametric forecasting of low-dimensional dynamical systems. *Phys. Rev. E.*, 91:032915, 2015.
- [8] T. Berry and J. Harlim. Variable bandwidth diffusion kernels. *Appl. Comput. Harmon. Anal.*, 2015. In press.
- [9] T. Berry and J. Harlim. Semiparametric forecasting and filtering: Correcting low-dimensional model error in parametric models. *J. Comput. Phys.*, 308, 2016. In press.
- [10] T. Berry and T. Sauer. Local kernels and the geometric structure of data. *Appl. Comput. Harmon. Anal.*, 2015. In press.
- [11] E. Bodenschatz, S. P. Malinowski, R. A. Saw, and F. Stratmann. Can we understand clouds without turbulence? *Science*, 327:970–971, 2010.
- [12] G. Branstator, H. Teng, G. A. Meehl, M. Kimoto, J. R. Knight, M. Latif, and A. Mosati. Systematic estimates of initial-value decadal predictability for six AOGCMs. *J. Climate*, pages 1827–1846, 2012.
- [13] D. S. Broomhead and G. P. King. Extracting qualitative dynamics from experimental data. *Phys. D*, 20(2–3):217–236, 1986.
- [14] M. Bushuk, D. Giannakis, and A. J. Majda. Reemergence mechanisms for North Pacific sea ice revealed through nonlinear Laplacian spectral analysis. *J. Climate*, 27(16):6265–6287, 2014.
- [15] N. Chen, A. J. Majda, and D. Giannakis. Predicting the cloud patterns of the Madden-Julian Oscillation through a low-order nonlinear stochastic model. *Geophysical Research Letter*, 41(15):5612–5619, 2014.
- [16] R. R. Coifman and M. J. Hirn. Bi-stochastic kernels via asymmetric affinity functions. *Appl. Comput. Harmon. Anal.*, 35:177–180, 2013.

- [17] R. R. Coifman and S. Lafon. Diffusion maps. *Appl. Comput. Harmon. Anal.*, 21:5–30, 2006.
- [18] R. R. Coifman and S. Lafon. Geometric harmonics: A novel tool for multiscale out-of-sample extension of empirical functions. *Appl. Comput. Harmon. Anal.*, 21:31–52, 2006.
- [19] W. D. Collins, C. M. Bitz, M. L. Blackmon, G. B. Bonan, C. S. Bretherton, J. A. Carton, P. Chang, S. C. Doney, J. Hack, T. B. Henderson, J. T. Kiehl, W. G. Large, D. S. McKenna, B. D. Santer, and R. D. Smith. The community climate system model version 3 (CCSM3). *J. Climate*, 19:2122–2143, 2006.
- [20] D. Comeau, Z. Zhao, D. Giannakis, and A. J. Majda. Data-driven prediction strategies for low-dimensional patterns of North Pacific climate variability, in revision.
- [21] S. Corti, A. Weisheimer, T. N. Palmer, F. J. Doblas-Reyes, and L. Magnusson. Reliability of decadal predictions. *Geophys. Res. Lett.*, 39:L21712, 2012.
- [22] N. Cressie and C. K. Wikle. *Statistics for Spatio-Temporal Data*. Wiley, Hoboken, 2011.
- [23] D. T. Crommelin and A. J. Majda. Strategies for model reduction: Comparing different optimal bases. *J. Atmos. Sci.*, 61:2206–2217, 2004.
- [24] Q. Deng, B. Khouider, and A. J. Majda. The MJO in a coarse-resolution GCM with a stochastic multicloud parameterization. *J. Atmos. Sci.*, 2014. In press.
- [25] V. P. Dymnikov and A. N. Filatov. *Mathematics of Climate Modeling*. Birkhäuser, Boston, 1997.
- [26] T. Eisner, B. Farkas, M. Haase, and R. Nagel. *Operator Theoretic Aspects of Ergodic Theory*, volume 272 of *Graduate Texts in Mathematics*. Springer, 2015.
- [27] J. Farmer and J. Sidorowich. Predicting chaotic time series. *Phys. Rev. Lett.*, 59(8):845–848, 1987.
- [28] A. Fernández, N. Rabin, D. Fishelov, and J. R. Dorronsoro. Auto-adaptive Laplacian pyramids for high-dimensional data analysis. arxiv.org/abs/1311.6594, 2014.
- [29] B. Fox-Kemper, S. Bachman, B. Pearson, and S. Reckinger. Principles and advances in subgrid modelling for eddy-rich simulations. *CLIVAR Exchanges*, 65(2):42–46, 2014.
- [30] D. Giannakis. Dynamics-adapted cone kernels. *SIAM J. Appl. Dyn. Sys.*, 14(2):556–608, 2015.
- [31] D. Giannakis and A. J. Majda. Comparing low-frequency and intermittent variability in comprehensive climate models through nonlinear laplacian spectral analysis. *Geophys. Res. Lett.*, 39:L10710, 2012.
- [32] D. Giannakis and A. J. Majda. Nonlinear Laplacian spectral analysis for time series with intermittency and low-frequency variability. *Proc. Natl. Acad. Sci.*, 109(7):2222–2227, 2012.
- [33] D. Giannakis and A. J. Majda. Nonlinear Laplacian spectral analysis: Capturing intermittent and low-frequency spatiotemporal patterns in high-dimensional data. *Stat. Anal. Data Min.*, 6(3):180–194, 2013.
- [34] D. Giannakis and A. J. Majda. Data-driven methods for dynamical systems: Quantifying predictability and extracting spatiotemporal patterns. In R. Melnik, editor, *Mathematical and Computational Modeling: With Applications in Engineering and the Natural and Social Sciences*, page 288. Wiley, Hoboken, 2014.
- [35] M. Hein, J.-Y. Audibert, and U. von Luxburg. From graphs to manifolds—weak and strong pointwise consistency of graph Laplacians. In *Learning Theory*, volume 3559 of *Lecture Notes in Computer Science*, pages 470–485. Springer, Berlin, 2005.
- [36] H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, Cambridge, 2004.
- [37] A. Katok and Hasselblatt B. *Introduction to the Modern Theory of Dynamical Systems*. Cambridge University Press, Cambridge, 1997.
- [38] U. Krengel. *Ergodic Theorems*, volume 6 of *De Gruyter Studies in Mathematics*. Walter de Gruyter, Berlin, 1985.
- [39] D. Kugiumtzis, O. C. Lingjaerde, and N. Christophersen. Regularized local linear prediction of chaotic time series. *Physica D: Nonlinear Phenomena*, 112(3):344–360, 1998.
- [40] W. K. M. Lau and D. E. Waliser. *Intraseasonal Variability in the Atmosphere–Ocean Climate System*. Springer-Verlag, Berlin, 2011.
- [41] E. N. Lorenz. Atmospheric predictability as revealed by naturally occurring analogues. *J. Atmos. Sci.*,

- 26:636 – 646, 1969.
- [42] A. J. Majda and J. Harlim. Physics constrained nonlinear regression models for time series. *Nonlinearity*, 26(1):201–217, 2013.
 - [43] N. J. Mantua and S. R. Hare. The Pacific Decadal Oscillation. *J. Oceanogr.*, 58:35–44, 2002.
 - [44] G. A. Meehl et al. Decadal prediction: Can it be skillful? *Bull. Amer. Meteor. Soc.*, 90(10):1467–1485, 2009.
 - [45] M. W. Moncrieff. Collaborative research at the intersection of weather and climate. *Bull. Amer. Meteor. Soc.*, 56(3):204–211, 2007.
 - [46] M. Newman, G. P. Compo, and M. A. Alexander. ENSO-forced variability of the Pacific Decadal Oscillation. *J. Climate*, 16:3853–3857, 2003.
 - [47] E. J. Nyström. Über die praktische Auflösung von Integralgleichungen mit Anwendungen auf Randwertaufgaben. *Acta Mathematica*, pages 185–204, 1930.
 - [48] T. N. Palmer and R. Hagedorn, editors. *Predictability of Weather and Climate*. Cambridge University Press, Cambridge, 2006.
 - [49] Y. B. Pesin. Families of invariant manifolds corresponding to nonzero characteristic exponents. *Math*, 10:1261–1305, 1976.
 - [50] N. Rabin and R. R. Coifman. Heterogeneous datasets representation and learning using diffusion maps and Laplacian pyramids. In *12th SIAM International Conference on Data Mining*, 2012.
 - [51] T. Sauer. Time series prediction by using delay coordinate embedding. In *Time Series Prediction, Forecasting the Future and Understanding the Past*, pages 175–193, 1994.
 - [52] T. Sauer, J. A. Yorke, and M. Casdagli. Embedology. *J. Stat. Phys.*, 65(3–4):579–616, 1991.
 - [53] A. Singer. From graph to manifold Laplacian: The convergence rate. *J. Appl. Comput. Harmon. Anal.*, 21:128–134, 2006.
 - [54] A. Singer, R. Erban, I. G. Kevekidis, and R. R. Coifman. Detecting intrinsic slow variables in stochastic dynamical systems by anisotropic diffusion maps. *Proc. Natl. Acad. Sci.*, 106(38):16090–16095, 2009.
 - [55] G. Sugihara and R. M. May. Nonlinear forecasting as a way of distinguishing chaos from measurement error in time series. *Nature*, 344(6268):734–741, 1990.
 - [56] E. Székely, D. Giannakis, and A. J. Majda. Extraction and predictability of coherent intraseasonal signals in infrared brightness temperature data. *Climate Dyn.*, 2014. Submitted.
 - [57] F. Takens. Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence, Warwick 1980*, volume 898 of *Lecture Notes in Mathematics*, pages 366–381. Springer, Berlin, 1981.
 - [58] R. Talmon and R. R. Coifman. Empirical intrinsic geometry for nonlinear modeling and time series filtering. *Proc. Natl. Acad. Sci.*, 110(31):12535–12540, 2013.
 - [59] D. Ting, L. Huang, and M. Jordan. An analysis of the convergence of graph laplacians. In *Proceedings of the 27th International Conference on Machine Learning*, Haifa, Israel, 2010.
 - [60] Z. Toth. Long-range weather forecasting using an analog approach. *J. Climate*, 2:594 – 607, 1989.
 - [61] K. E. Trenberth. The definition of El Niño. *Bull. Amer. Meteor. Soc.*, 78:2771–2777, 1997.
 - [62] H. van den Dool, J. Huang, and Y. Fan. Performance and analysis of the constructed analogue method applied to U. S. soil moisture over 1981–2001. *J. Geophys. Res.*, 108(D16):8617, 2003.
 - [63] U. von Luxburg, O. Bousquet, and M. Belkin. On the convergence of spectral clustering on random samples: The normalized case. In *17th Annual Conference on Learning Theory, COLT 2004*, pages 457–471. Springer, 2004.
 - [64] M. Wheeler and K. M. Weickmann. Real-time monitoring and prediction of modes of coherent synoptic to intraseasonal tropical variability. *Mon. Wea. Rev.*, 129:2677–2694, 2001.
 - [65] P. K. Xavier and B. N. Goswami. An analog method for real-time forecasting of summer monsoon subseasonal variability. *Mon. Wea. Rev.*, 135:4149–4160, 2007.